

ERASING THE VARIABLE: EMPIRICAL FOREGROUND DISCOVERY FOR GLOBAL 21 CM SPECTRUM EXPERIMENTS

ERIC R. SWITZER^{1,2}, ADRIAN LIU^{3,4}

Submitted to ApJ

ABSTRACT

Spectral measurements of the 21 cm monopole background have the promise of revealing the bulk energetic properties and ionization state of our universe from $z \sim 6 - 30$. Synchrotron foregrounds are orders of magnitude larger than the cosmological signal, and are the principal challenge faced by these experiments. While synchrotron radiation is thought to be spectrally smooth and described by relatively few degrees of freedom, the instrumental response to bright foregrounds may be much more complex. To deal with such complexities, we develop an approach that discovers contaminated spectral modes using spatial fluctuations of the measured data. This approach exploits the fact that foregrounds vary across the sky while the signal does not. The discovered modes are projected out of each line-of-sight of a data cube. An angular weighting then optimizes the cosmological signal amplitude estimate by giving preference to lower-noise regions. Using this method, we show that it is essential for the passband to be stable to at least $\sim 10^{-4}$. In contrast, the constraints on the spectral smoothness of the absolute calibration are mainly aesthetic if one is able to take advantage of spatial information. To the extent it is understood, controlling polarization to intensity leakage at the $\sim 10^{-2}$ level will also be essential to rejecting Faraday rotation of the polarized synchrotron emission.

Subject headings: dark ages, reionization, first stars – methods: data analysis – methods: statistical

1. INTRODUCTION

One of the richest yet least understood narratives in cosmology is the formation of the complex structure that we see today out of the simple initial conditions implied by the cosmic microwave background (CMB). The first luminous objects are thought to have formed at $z \sim 20 - 30$ through collapse in $10^6 - 10^8 M_\odot$ halos (Barkana & Loeb 2001; Bromm 2013). The radiation from these objects heated and then reionized the intergalactic medium (IGM). There are several sources of complementary information about the evolution of ionization in this epoch.

The CMB temperature anisotropy is damped by the total Thomson depth to free electrons. The Planck collaboration has used this effect, combined with a constraint on the scalar amplitude from gravitational lensing, to measure the optical depth through reionization (Planck Collaboration et al. 2013). In addition, Thomson scattering through the reionization epoch generates a unique polarization signature on large angular scales. WMAP has measured the total optical depth using this polarization signature (Bennett et al. 2013). These are integral constraints on the free electron abundance, and can be translated into a central reionization redshift of 10.6 ± 1.1 (Bennett et al. 2013).

Once the IGM is highly ionized, it is transparent to Lyman- α photons. Absorption measurements along sight lines to high-redshift quasars indicate that reionization must have ended by $z < 6$ (Fan et al. 2006). Absorption saturates at low abundance, so these should be taken as bounds on the end of reionization, which could still have been largely complete at redshifts higher than 6.

Recently, two methods have been developed to place much more direct bounds on the duration of reionization. The ionization process is thought to be spatially patchy, as local sources of radiation blow ionized bubbles that coalesce into the fully-reionized IGM. CMB photons scattering from this patchy screen produce an additional kinetic Sunyaev-Zel'dovich anisotropy appearing most clearly at $\ell > 3000$, where the primary CMB is negligible (Knox et al. 1998). Upper limits on this effect translate into a model-dependent upper bound on the duration of reionization (Zahn et al. 2012), and hold the promise of direct detection of patchy structure in the near future.

The patchy structure of reionization can also be observed directly in three dimensions using emission of neutral hydrogen through its 21 cm line (Furlanetto et al. 2006). Recent bounds from GMRT (Paciga et al. 2013), MWA (Dillon et al. 2014) and PAPER (Parsons et al. 2013; Pober et al. 2013) are marching down to the expected level of fluctuations, in parallel to efforts at LOFAR (van Haarlem et al. 2013). An alternative to measuring the 21 cm anisotropy is to measure the signal of its global emission (or absorption at earlier times) (Pritchard & Loeb 2008, 2010), which reveals the bulk energetic properties and ionization state of the universe during reionization and preceding epochs when the first luminous structures were forming. Global 21 cm experiments include EDGES (Bowman & Rogers 2010) and SCI-HI (Voytek et al. 2014) (which have both reported bounds), LEDA⁵ and the proposed DARE mission (Burns et al. 2012).

The frequencies of interest in these global studies are $\sim 50 - 200$ MHz, and fiducial theoretical models suggest a maximum contrast of ~ 100 mK relative to the synchrotron emission of the galaxy, which can vary $\sim 10^2 - 10^5$ K across the sky and frequency range. Astrophysical synchrotron emission is thought to be fully described by a handful of spectrally smooth functions that can be distinguished from the variation

¹ NASA Goddard Space Flight Center, Greenbelt, MD 20771

² Canadian Institute for Theoretical Astrophysics, University of Toronto, 60 St. George St., Toronto, Ontario, M5S 3H8

³ Department of Astronomy, UC Berkeley, Berkeley, CA 94720, USA

⁴ Berkeley Center for Cosmological Physics, UC Berkeley, Berkeley, CA 94720, USA
eric.r.switzer@nasa.gov

⁵ <http://www.cfa.harvard.edu/LEDA/science.html>

of the global reionization signal.

Extremely bright foregrounds make the measurement susceptible to instrumental systematics. For example, if an instrument with a 1% ripple in one spectral band observed a 500 K power law, subtraction of a pure power law would leave a 5 K residual, significantly larger than the signal. Through careful instrumental design, the level of instrumental systematics may be controlled, but generally cannot be nulled entirely.

Here, we develop methods that can be used to constrain the cosmological signal in these heavily contaminated data. Following the monopole nature of the signal, experiments to-date have mapped the sky with very large beams (Bowman & Rogers 2010; Voytek et al. 2014). However, a unique trait of the foregrounds is that they vary across the sky, while the signal is constant. This regime is the opposite of the situation normally found in analysis of small signals, where a modulated signal is pulled out of foregrounds. Liu et al. (2013) (henceforth L13) proposed that experiments seeking to measure the global signal should also resolve the sky with an instrumental beam. This allows selective weighting against regions of high contamination, and allows for the use of angular correlation information to reject foregrounds. The additional spatial resolution yields higher fidelity recovery of the cosmological 21 cm spectrum.

Here, we extend the ideas in L13 to a method that uses the spatial fluctuations of foregrounds in the data to discover contaminated spectral modes. A similar idea has been employed successfully in 21 cm large scale structure measurements (Switzer et al. 2013; Masui et al. 2013; Chang et al. 2010) at $z \sim 1$, and has been suggested for cleaning ionospheric contamination (Vedantham et al. 2014).

Discovery of foreground spectral modes in the measured data makes the method more robust to assumptions about the foreground contamination. For example, now if the instrumental passband has a 1% ripple, the largest foreground mode discovered in our foreground cleaning method will also self-consistently exhibit this ripple. Generally, instrumental systematics take relatively clean and smooth functions of frequency from synchrotron emission and convert them into more complex structure that requires additional spectral functions to describe. We argue that the primary goal in instrumental design should be to prevent proliferation of bright, new foreground modes in the data. Each new foreground degree of freedom produced by instrumental response to foregrounds results in more signal loss and makes discovery of the signal more ambiguous.

The methods described here of 1) using spatial variation to discover spectral foreground modes, which can then be projected out, and 2) downweighting known spatial areas of high contamination (the galaxy) provide the strongest methods for recovering the global 21 cm signal in the absence of additional prior information about the foregrounds or instrumental response. While the algorithm of mode subtraction and angular weighting is intuitive, we develop it from the ground up to expose several implicit choices and possible pitfalls.

Recently, Bernardi et al. (2014) argued that a dipole gain pattern can be calibrated in an interferometric array. However, additional variations in spectral response due to factors such as the analog to digital converter, reflection and signal loss after the antenna were not included. Our goal here is to understand how data analysis can be made more robust to this class of instrumental response (or any other source of foreground covariance), or alternately how tightly certain instru-

mental tolerances must be constrained.

In Sec. 2 we review the basic properties of the global signal and describe our foreground model. Sec. 3 builds up the estimator for joint foreground and signal estimation using a simplified model of spectra along independent sightlines. Sec. 4 considers implications of this model for passband calibration. Sec. 5 develops spatial weights, and Sec. 6 combines the estimators with spatial and spectral weights. Sec. 7 describes a number of considerations for using the methods developed here and for global 21 cm signal estimation in general. We discuss telescope beamwidth, the foreground monopole, how aggressively foreground modes should be removed, and susceptibility to Faraday rotation. We also consider mode removal of the pre-reionization absorption feature and extensions of the simple template amplitude constraint considered throughout. We summarize our conclusions in Sec. 8.

2. THE GLOBAL SIGNAL AND FOREGROUND MODELS

From radiative transport arguments, the brightness temperature of 21 cm radiation is

$$T_b \approx 27(1 - \bar{x}_i) \left(\frac{T_S - T_{CMB}}{T_S} \right) \left(\frac{1+z}{10} \right)^{1/2} \text{ mK}, \quad (1)$$

where \bar{x}_i is the mean ionization fraction, T_S is the spin temperature of the hyperfine transition, and T_{CMB} is the CMB temperature. The basic physics of the spin-temperature coupling is well-understood (Pritchard & Loeb 2010), but the detailed astrophysical processes that determine the coupling strength and the gas temperature are still conjectural.

For $z > 200$, the universe is dense enough that electron collisional interactions drive the spin temperature to the gas temperature, which is cooling faster than the CMB. This produces absorption. By $z \sim 30$, the universe is sufficiently rarefied that the spin temperature is better coupled to the CMB bath, and absorption is expected to subside. Once the first luminous objects form, these produce radiation that drives the Wouthuysen-Field (Wouthuysen 1952; Field 1958) coupling of T_S again to the gas temperature, leading to a second absorption feature. Then, X-ray heating of the gas can drive the spin temperature above the CMB temperature, leading to emission. As these luminous processes proceed and increase, they also ionize the IGM, which causes the signal to disappear as $1 - \bar{x}_i \rightarrow 0$ (Pritchard & Loeb 2010).

We will find it convenient to have a reionization model with a small number of parameters, rather than a full spectrum. For concreteness, we will spend most of the paper focusing on the evolution of the brightness temperature as the universe is re-ionized, rather than the absorption dip as heating begins. However, our methods are applicable at all redshifts, and in Sec. 7.6 we will briefly discuss the pre-reionization dip.

At the beginning of the reionization epoch, it is thought that the spin temperature is strongly coupled to the gas, and that the gas is heated, driving 21 cm emission. As the ionization proceeds, this emission dies away. A simple way of parameterizing this is (Pritchard & Loeb 2010)

$$\bar{x}_i(z) = \frac{1}{2} \left[1 + \tanh \left(\frac{z_r - z}{\Delta z} \right) \right], \quad (2)$$

so that the brightness temperature scales as (Fig. 1)

$$T_b(z) = \frac{27}{2} \left(\frac{1+z}{10} \right)^{1/2} \left[1 + \tanh \left(\frac{z_r - z}{\Delta z} \right) \right] \text{ mK} \quad (3)$$




Figure 1. Brightness temperature of the global 21 cm emission in a simplified two-parameter model of reionization. The central redshift is taken from the WMAP9 constraint at $z = 10.6$, and several representative values for the duration of reionization are shown.




Figure 2. Foreground model at 80 MHz, displayed as $\log(T)$.

where we have assumed that $T_S \gg T_{CMB}$.

Our model for the diffuse intensity (Stokes-I) of the foregrounds is based on an extended version of the Global Sky Model (GSM) of de Oliveira-Costa et al. (2008) that is developed in L13. The original GSM used a principal component analysis to extrapolate and interpolate between previous galactic surveys at a wide range of frequencies. It is based on 3 spectral components and is accurate down to $\approx 10^\circ$. L13 extends this model by adding additional (less smooth) spectral eigenmodes, as well as a population of power-law sources with randomly distributed spectral indices, drawn from the dN/dS brightness distribution of Di Matteo et al. (2002). The goal of adding these components is to boost the rank of the foreground covariance to reflect spectral fluctuations that can be expected on the real sky. These cannot be probed in the original rank-3 GSM. An example map of our foreground model at 80 MHz is shown in Fig. 2.

In Sec. 7.7, we will additionally consider the impact of Faraday rotation and polarized foregrounds. As this is somewhat peripheral to the development of our data analysis methods in the next few sections, we will defer our discussion of our foreground polarization model until then.

3. GLOBAL 21 CM OBSERVATIONS ALONG INDEPENDENT LINES OF SIGHT

A global 21 cm experiment that maps the sky will have a finite number of resolution elements across its field, defined by the beam. In each one of these resolution elements, it observes the spectrum of incoming radiation. A reasonable approximation is to split the survey into N_θ spectra of length N_ν along all the angular resolution elements. Let \mathbf{y}_i be that vector, where i indexes the sight line (1 to N_θ) and the length of the vector is N_ν . In this section, we will build up infrastructure and intuition using this simplified case of N_θ independent and identically-distributed sight lines. Sec. 6 extends this to a method that treats realistic foregrounds with proper angular correlations. Because angular correlations become largely irrelevant when considering wide-beam experiments with essentially no angular sensitivity, much of the intuition developed here is directly transferable to the analysis of such experiments.

In writing an estimator, there are many choices for aspects of the global signal that could be estimated. The goal could be to constrain 1) an arbitrary spectrum of the cosmological 21 cm evolution, 2) some modes of its variation that relate to physical parameters, 3) some amplitudes that are based on external information about the expected redshift of reionization, 4) the amplitude of a template of a provisional global 21 cm signal. For our purposes here, it is simplest to develop estimators for the amplitude of an assumed 21 cm signal template. Initially, experiments will simply be stepping down in bounds and seeking some evidence of $z \sim 20$ absorption and heating or a $z \sim 11$ reionization signal—an amplitude constraint could provide the simplest, clear indication of a cosmological signal. Sec. 7.5 considers other estimation regimes in light of the methods developed here.

Let the assumed template of the global signal be a vector \mathbf{x} of length N_ν and normalized to have a maximum of 1. Multiply by some amplitude α to get the 21 cm signal. Then the observed spectrum is the sum of signal, thermal noise $\mathbf{n}_{\text{inst},i}$ and foregrounds $\mathbf{n}_{\text{fg},i}$

$$\mathbf{y}_i = \alpha \mathbf{x} + \mathbf{n}_{\text{fg},i} + \mathbf{n}_{\text{inst},i}. \quad (4)$$

Initially we will assume that the foregrounds and noise are identically and independently normally distributed along each line of sight i , $\mathbf{n}_{\text{fg},i} \sim N(0, \Sigma_{\text{fg}})$ and $\mathbf{n}_{\text{inst},i} \sim N(0, \Sigma_{\text{inst}})$. In reality, the foreground field is strongly correlated in angle and non-Gaussian. These issues will be examined in subsequent sections. In contrast, thermal noise is uncorrelated between sight lines and normally distributed: $\mathbf{n}_{\text{inst},i} \sim N(0, \Sigma_{\text{inst}})$ will remain an excellent approximation. Throughout, the $N_\nu \times N_\nu$ matrix $\Sigma \equiv \Sigma_{\text{inst}} + \Sigma_{\text{fg}}$ will refer to the total (ν, ν') covariance.

3.1. Known covariance

If the foregrounds and thermal noise are drawn from a total, known covariance Σ , then the maximum likelihood estimate for the template amplitude is

$$\hat{\alpha}_{\text{ML}} = (\mathbf{x}^T \Sigma^{-1} \mathbf{x})^{-1} \mathbf{x}^T \Sigma^{-1} \bar{\mathbf{y}} \quad (5)$$

where

$$\bar{\mathbf{y}} = N_\theta^{-1} \sum_{i=1}^{N_\theta} \mathbf{y}_i \quad (6)$$

is the mean spectrum along all lines of sight.

This can be understood as “deweighting” the foregrounds ($\Sigma^{-1}\bar{y}$), projecting onto the signal template ($x^T \Sigma^{-1}\bar{y}$) and then applying a normalization to account for the weights. The estimated amplitude is normally distributed and its error is

$$\text{Var}(\hat{\alpha}_{\text{ML}})|_{\Sigma} = (x^T \Sigma^{-1} x)^{-1}. \quad (7)$$

To convert this into a more intuitive cleaning process, take the eigendecomposition of the covariance $\Sigma = V \Lambda V^T$. The cleaning operation $\Sigma^{-1}\bar{y}$ is then $V \Lambda^{-1} V^T \bar{y}$. Here V^T projects the data onto a basis where the contaminated modes v_j (the j ’th row of V) are orthonormal. Those modes are weighted against $\lambda_j^{-1} = [\Lambda^{-1}]_{jj}$, where the highest variance, most contaminated modes are down-weighted. Then, V projects this down-weighted basis back onto the spectral basis.

In the usual assumptions, the foregrounds are not full rank, even if the spectral covariance Σ —which contains instrumental noise—is full rank. Often, the foregrounds are instead described as some set of contaminated modes v_j where j ranges from 1 to the number of contaminated modes N_{fg} , or

$$\mathbf{n}_{\text{fg},i} = \sum_j^{N_{\text{fg}}} a_{i,j} \mathbf{v}_j. \quad (8)$$

If the foreground covariance is described by only N_{fg} highly contaminated modes, a robust cleaning method is to null those modes entirely, setting $\lambda \rightarrow \infty$ artificially (Liu & Tegmark 2012). This is equivalent to fitting and subtracting those spectral modes from each line of sight, as

$$\bar{y}_{\text{clean}} = \sum_j^{N_{\text{fg}}} (1 - \mathbf{v}_j \mathbf{v}_j^T) \bar{y}, \quad (9)$$

where the \mathbf{v}_i are normalized so that $\mathbf{v}_i^T \mathbf{v}_i = 1$.

Here we have assumed that the foreground spectral functions are known in advance and can be projected out. This is essentially the same as arguments for subtracting polynomials or power laws along the line of sight. A crucial difference is that there is no assumption of smoothness—if the foregrounds were known to have a particular spectral shape, that could be represented in the vectors \mathbf{v}_i . The essential property of the foregrounds that allows for their removal is not that they are smooth, but that they are described by few functions. In the limit that the number of orthogonal functions removed approaches N_{ν} , all of the signal is removed because $\{\mathbf{v}_i, i = 1 \dots N_{\nu}\}$ spans the space, assuming orthogonality of the \mathbf{v}_i . The rank of bright foregrounds is the primary determinant of cosmological signal that can be extracted.

Subtraction of bright foregrounds has proven to be significantly more challenging in practice due to instrumental effects (Switzer et al. 2013). Fig. 3 shows an example of 0.1% calibration for a power law foreground spectrum similar in amplitude to the one reported by Voytek et al. (2014), and scaled to the slightly different frequencies of interest here. This represents the case where the \mathbf{v}_i mode removed assumes a pure power law, but the actual measurement reflects the instrument’s response to \mathbf{v}_i . Residuals are considerably larger than the signal.

One approach to treating the passband is to control it through instrumental design, and to measure it very precisely prior to the experiment. In principle, this could be achieved by long integrations on a calibration reference that is known




Figure 3. Impact of calibration uncertainty. *Top:* a power-law foreground spectrum similar to the one observed in Voytek et al. (2014), multiplied by a 0.1% calibration error. *Bottom:* residuals when a smooth power-law spectrum is subtracted from data with a calibration error. For contrast, the global signal at these redshifts is expected to be ~ 30 mK. Also note that calibration error will not integrate down thermally.

to be smooth. This carries the challenge of calibrating on a source other than the sky. An alternative is to recognize that the foregrounds are a bright spectral reference, viewed through the same survey program as the primary science. One calibration scheme would then consist of integrating down on bright synchrotron emission and renormalizing the spectrum under the assumption that the emission is spectrally smooth (Voytek et al. 2014).

An alternative to calibrating the instrument to meet the smooth mode functions $\{\mathbf{v}_i\}$ is to adjust the mode functions to reflect imperfections in the instrument. In the next section we develop a method for discovering foreground spectral mode functions in the data based on their *spatial* variations. The central idea is that the cosmological signal is a monopole spectrum, so the spectrum of anything that fluctuates across the sky must reflect some foreground modes (hence the title “Erasing the variable”). If there are monopole foreground components, they are formally indistinguishable from the cosmological 21 cm monopole in this picture (in the absence of prior information).

Measurement of $\{\mathbf{v}_i\}$ modes within the data may also discover unanticipated or unconstrained instrumental response systematics. Examples include 1) a passband that varies with time or pointing, which causes a response to synchrotron radiation that varies across the sky, 2) polarization leakage, which can cause Faraday (spectral) rotation that varies across the sky, 3) terrestrial interference that varies with instrument pointing, and 4) un-modeled frequency dependence of the beam. Even if the foregrounds were intrinsically a simple rank-1 spectral function, the instrument response will tend to proliferate the rank of the foreground covariance. This limits the ultimate sensitivity to the global signal. While many of these sources of systematics may be tightly controlled by the instrument’s construction, residual imperfections will not be known in advance and could be discovered by this method.

We will assume throughout that the true synchrotron foregrounds are described by a small number of modes $N_{\text{fg,true}}$. The instrument observes these bright foregrounds and records spectral data that may need some $N_{\text{fg}} \geq N_{\text{fg,true}}$ functions to describe the contamination.

3.2. Covariance from the measurements

This section considers the joint determination of the global signal amplitude α and the contaminated foreground modes. The statistical methods described and extended here were developed by Rao (1967). The central idea in these methods is to begin with a poor, contaminated estimator for the 21 cm signal amplitude. This initial estimator can then be cleaned up by projecting out correlations with the foregrounds.

To reflect the reality that often one does not know the covariance *a priori*, first consider a simple estimator that does not use any information from the spectral covariance. This will serve as our starting point for constructing an estimator that simultaneously estimates the spectral covariance from the data. Again, let \mathbf{x} be the spectral template of the cosmological signal that we would like to constrain. The simplest (and arguably worst) estimator we could develop for the amplitude is

$$T(\parallel \mathbf{x}) = (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \bar{\mathbf{y}}, \quad (10)$$

which takes the dot product of the 21 cm template and observed data ($\mathbf{x}^T \bar{\mathbf{y}}$) and normalizes through $(\mathbf{x}^T \mathbf{x})^{-1}$. We refer to this as $T(\parallel \mathbf{x})$ because it estimates the amplitude of the component of $\bar{\mathbf{y}}$ parallel to the signal \mathbf{x} . This is exactly the estimator that would be used if there were no frequency correlations in the noise. What makes this a very poor estimator is that it is strongly correlated with the foreground modes, and will therefore contain strong foreground residuals. In contrast, Eq. (5) in the previous section used the known frequency covariance to de-weight the foreground contamination.

We can begin to improve upon our simple estimator by constructing spectral modes that contain no cosmological signal by design. Using a Gram-Schmidt process, we can make a set of vectors \mathbf{z}_i that span the space orthogonal to the theoretical template of the cosmological signal \mathbf{x} , so that $\mathbf{z}_i^T \mathbf{x} = 0$. There are $N_\nu - 1$ such vectors, and the choice of these vectors is completely arbitrary at this point. Any spectrum can be represented as a sum

$$\mathbf{y} = \alpha \mathbf{x} + \sum_i \beta_i \mathbf{z}_i \quad (11)$$

because $\{\mathbf{x}, \mathbf{z}_i\}$ span the N_ν -dimensional space. We can pack these vectors \mathbf{z}_i into a matrix \mathbf{Z} and write an estimator for all the spectral information orthogonal to \mathbf{x} , $T(\perp \mathbf{x}) = \mathbf{Z}^T \bar{\mathbf{y}}$. This represents all the components of $\bar{\mathbf{y}}$ orthonormal to the signal, and is dominated by foregrounds. (Note that we could choose to normalize this for general vectors \mathbf{z} and find $(\mathbf{Z}^T \mathbf{Z})^{-1} \mathbf{Z}^T \bar{\mathbf{y}}$. This would make the equations more complex, but not modify the estimator for the cosmological quantity.)

The covariance between $T(\parallel \mathbf{x})$ and $T(\perp \mathbf{x})$ is

$$\begin{aligned} \text{Cov}_{\parallel, \perp} &= \begin{pmatrix} \mathbf{C}_{\parallel, \parallel} & \mathbf{C}_{\parallel, \perp} \\ \mathbf{C}_{\perp, \parallel} & \mathbf{C}_{\perp, \perp} \end{pmatrix} \\ &= \begin{pmatrix} (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \Sigma \mathbf{x} (\mathbf{x}^T \mathbf{x})^{-1} & (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \Sigma \mathbf{Z} \\ \mathbf{Z}^T \Sigma \mathbf{x} (\mathbf{x}^T \mathbf{x})^{-1} & \mathbf{Z}^T \Sigma \mathbf{Z} \end{pmatrix}. \end{aligned} \quad (12)$$

The off-diagonal terms $(\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \Sigma \mathbf{Z}$ represent correlations between $T(\parallel \mathbf{x})$ and $T(\perp \mathbf{x})$ caused by the (ν, ν') covariance of foregrounds in Σ . If the noise terms were only thermal fluctuations, $\Sigma = \sigma_{\text{inst}}^2 \mathbf{1}$, and by the construction $\mathbf{x}^T \mathbf{Z} = 0$, the off-diagonal terms vanish. In other words, an estimator $T(\perp \mathbf{x})$ that is known to contain only contaminants

(and no cosmological signal) is correlated with our estimate of the cosmological signal, so the latter must be contaminated.

With knowledge of spectral correlations, we can form an improved, adjusted estimator

$$\hat{\alpha} = T(\parallel \mathbf{x}) - \mathbf{C}_{\parallel, \perp} \mathbf{C}_{\perp, \perp}^{-1} T(\perp \mathbf{x}) \quad (13)$$

that projects foreground frequency correlations out of the $T(\parallel \mathbf{x})$ estimator. Intuitively, this estimator instructs us to make an estimate $T(\perp \mathbf{x})$ of the portion the measurement that is known to contain only foregrounds. Off-diagonal elements of the covariance matrix then allow the level of foreground leakage (into our estimate $T(\parallel \mathbf{x})$ of the cosmological signal) to be predicted and subtracted off. A similar technique was used recently in a power spectrum analysis of PAPER data (Parsons et al. 2013).

So far, we have assumed that Σ is known, but the same covariance adjustment can be performed with respect to an estimated $\hat{\Sigma}$. Without perfect knowledge of Σ , the error bars grow but the estimator remains unbiased. The (ν, ν') covariance can be estimated empirically using

$$\hat{\Sigma} = (N_\theta - 1)^{-1} \sum_{i=1}^{N_\theta} (\mathbf{y}_i - \bar{\mathbf{y}})(\mathbf{y}_i - \bar{\mathbf{y}})^T \quad (14)$$

Writing out Eq. (13),

$$\hat{\alpha} = (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T (1 - \mathbf{\Pi}) \bar{\mathbf{y}} \quad (15)$$

where $\mathbf{\Pi} = \hat{\Sigma} \mathbf{Z} (\mathbf{Z}^T \hat{\Sigma} \mathbf{Z})^{-1} \mathbf{Z}^T$. Remarkably, under our current assumptions where \mathbf{Z} spans the rest of the N_ν dimensions, this estimator is equivalent to (see Appendix A)

$$\hat{\alpha} = (\mathbf{x}^T \hat{\Sigma}^{-1} \mathbf{x})^{-1} \mathbf{x}^T \hat{\Sigma}^{-1} \bar{\mathbf{y}}, \quad (16)$$

which is precisely the same as Eq. (5), except with the estimated spectral covariance $\hat{\Sigma}$ replacing the known covariance Σ .

Gleser & Olkin (1972) develop the same result through a completely independent method. They write down the joint likelihood of the covariance and mean (cosmological signal), where the covariance is Wishart-distributed and the mean is normally distributed. They then maximize the likelihood and find the same result, showing that the choice $\Sigma \rightarrow \hat{\Sigma}$ coincides with the maximum likelihood.

In summary, the proposed procedure is to

- Find the sample mean and (ν, ν') covariance across observed lines of sight.
- Invert the measured covariance and use it to down-weight contaminated modes in the data.
- Find the inner product of the cleaned data and the signal template, then normalize.

While the maximum of the likelihood has the same form as the case where the covariance is known, the distribution of the estimated 21 cm global signal template amplitude is no longer Gaussian, and is generally broader. These changes are due to the fact that the data are also used for the covariance estimation, which uses up degrees of freedom in determining the foreground modes. Rao (1967) shows that

$$\text{Var}(\hat{\alpha}) = \frac{N_\theta - 1}{N_\theta - 1 - r} [\mathbf{C}_{\parallel, \parallel} - \mathbf{C}_{\parallel, \perp} \mathbf{C}_{\perp, \perp}^{-1} \mathbf{C}_{\perp, \parallel}], \quad (17)$$

where r is the rank of the additional degrees of freedom that are estimated. In our case here, $r = N_\nu - 1$ (\mathbf{Z} spans the rest of the spectral space) so that

$$\text{Var}(\hat{\alpha}) = \left(1 + \frac{N_\nu - 1}{N_\theta - N_\nu}\right) (\mathbf{x}^T \hat{\Sigma}^{-1} \mathbf{x})^{-1}, \quad (18)$$

where we have also plugged in the explicit forms for $\mathbf{C}_{\parallel,\parallel}$, $\mathbf{C}_{\parallel,\perp}$, $\mathbf{C}_{\perp,\perp}$, and $\mathbf{C}_{\perp,\parallel}$ from Eq. (12). We see that the covariance is enhanced relative to the case of perfect foreground covariance knowledge [Eq. (7)] by a factor related to the number of frequencies and independent lines of sight. The error diverges when N_θ approaches N_ν from above. The rank $r(\hat{\Sigma}) \leq \min(N_\theta, N_\nu)$, and for $N_\theta \gg N_\nu$, the covariance is well-measured. This suggests that the optimal limit is to have many more resolution elements than spectral bins.

There is little instrumental limitation on the number of spectral bins N_ν over the bands observed. In contrast, there is a hard limit on the number of “independent” lines of sight that can be observed, N_θ . At the frequencies of interest for the global 21 cm signal, the beam size tends to be large by diffraction. Angular correlations and noise generally limit the number of independent samples of the spectrum (see Sec. 7.1).

The requirement $N_\theta > N_\nu$ stems from the fact that $\hat{\Sigma}$ is trying to estimate a full N_ν rank covariance. In general, the sensitivity to α should not depend on the number of spectral bins in the survey, once they become fine enough to resolve the signal. The solution is to instead estimate some number N_{fg} of contaminated spectral functions, less than the number of spectral bins. Then the estimation becomes independent of the N_ν , so long as $N_\nu > N_{\text{fg}}$.

3.3. Empirically-determined foregrounds of limited rank

The key step in the covariance adjustment scheme of the previous section was the formation and application of $\mathbf{\Pi} = \hat{\Sigma} \mathbf{Z} (\mathbf{Z}^T \hat{\Sigma} \mathbf{Z})^{-1} \mathbf{Z}^T$, which modeled contaminated modes in the original “poor” estimator for later subtraction. The modes formed a basis \mathbf{Z} orthogonal to the signal \mathbf{x} . The specific choice of vectors in \mathbf{Z} was arbitrary so long as they were orthonormal (for simplicity) and spanned the spectral subspace orthogonal to the signal. For concreteness, we suggested forming this basis blindly using a Gram-Schmidt process. While formally a solution to our problem, this is not a particularly efficient way to implement our recipe in practice, for each of the resulting basis vectors will be a linear combination of noise and foreground modes. On the other hand, the foregrounds should be describable by a small number of modes (Liu & Tegmark 2012; Bernardi et al. 2014). A sensible alternative would therefore be to intelligently partition the basis \mathbf{Z} into two sub-bases, one that is of a relatively low rank N_{fg} containing the foregrounds, and another of rank $N_\nu - N_{\text{fg}} - 1$ consisting of the remaining noise-dominated modes. Computationally, this means that rather than estimating the full N_ν rank covariance from the data, we only need to estimate a subset of contaminated modes and their amplitudes. This limits the number of degrees of freedom r that need to be estimated from the data, which as we saw from Eq. (17), is crucial for keeping the final error bars small.

In the rank-restricted foreground approximation, the data along a line of sight are the cosmological signal plus some amplitudes β times foreground modes \mathbf{F} plus thermal noise, as

$$\mathbf{y} = \alpha \mathbf{x} + \mathbf{F} \beta + \mathbf{n}_t \Rightarrow \Sigma = \mathbf{F} \mathbf{\Gamma} \mathbf{F}^T + \sigma_{\text{inst}}^2 \mathbf{1}, \quad (19)$$

where we have taken the noise to be stationary for simplicity. The amplitude-ordered eigenvalue spectrum of the total covariance Σ in this case would show some large contaminant amplitudes followed by a noise floor set by σ .

Unlike \mathbf{Z} of the previous section, the foreground mode vectors in \mathbf{F} alone do not span the entire spectral subspace orthogonal to the cosmological signal. The remaining portion of the space is spanned by basis vectors that describe instrumental noise, under the assumptions of Eq. (19). These vectors can be formed, like before, using a Gram-Schmidt process (this time relative to the signal \mathbf{x} and foregrounds \mathbf{F}) and packed into a matrix \mathbf{G} . The observed spectrum \mathbf{y} can then be written as

$$\mathbf{y} = \alpha \mathbf{x} + \sum_{j=1}^{N_{\text{fg}}} \beta_j \mathbf{f}_j + \sum_{j=N_{\text{fg}}+1}^{N_\nu} \gamma_j \mathbf{g}_j, \quad (20)$$

which is to be contrasted with Eq. (11).

Again, we can form poor estimators for the amplitudes of the signal ($\hat{\alpha} = (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \bar{\mathbf{y}}$), foreground modes ($\hat{\beta} = \mathbf{F}^T \bar{\mathbf{y}}$) and thermal noise modes ($\hat{\gamma} = \mathbf{G}^T \bar{\mathbf{y}}$). To do so, however, we must first specify how the foreground modes \mathbf{F} are identified and defined, and in the following section we outline two different methods for this.

3.4. Two outlooks on separation of signal and foregrounds

The first method is closely related to the treatment that we have presented so far. In particular, we considered modes in the matrices \mathbf{Z} (full rank) or \mathbf{F} (finite rank) that represent components of the data that are orthogonal to the cosmological signal. There is, however, nothing preventing real foregrounds from being parallel to the signal, and generally \mathbf{n}_{fg} will be the sum of $\mathbf{n}_{\text{fg},\parallel\mathbf{x}}$ (foregrounds parallel to the cosmological signal) and $\mathbf{n}_{\text{fg},\perp\mathbf{x}}$ (foregrounds perpendicular to the signal). For example, if the signal and foregrounds share a slowly varying spectral component, it will fall in $\mathbf{n}_{\text{fg},\parallel\mathbf{x}}$.

An alternative is to form foreground spectral modes that best describe the covariance of the fluctuating terms on the sky, without regard for the cosmological signal. Then, the signal has a piece which is parallel to those foreground modes, and a piece which is perpendicular, or $\mathbf{x} = \mathbf{x}_{\parallel\mathbf{F}} + \mathbf{x}_{\perp\mathbf{F}}$.

In equations, these two methods are

Method 1: $\mathbf{y} = \mathbf{x} + \mathbf{n}_{\text{fg},\parallel\mathbf{x}} + \mathbf{n}_{\text{fg},\perp\mathbf{x}}$ delete: $\mathbf{n}_{\text{fg},\perp\mathbf{x}}$

Method 2: $\mathbf{y} = \mathbf{x}_{\perp\mathbf{F}} + \mathbf{x}_{\parallel\mathbf{F}} + \mathbf{n}_{\text{fg}}$ delete: $\mathbf{x}_{\parallel\mathbf{F}} + \mathbf{n}_{\text{fg}}$.

In the first method, we can develop a cleaning operation that removes/deweights $\mathbf{n}_{\text{fg},\perp\mathbf{x}}$, that is, components of the foreground variance orthogonal to the signal. This estimator does not touch the signal by design, so is guaranteed to have no cosmological signal loss. However, substantial foregrounds will remain in the estimated signal with this rather conservative approach. In the second method, we can develop cleaning that removes any foreground spectral modes that vary spatially on the sky, along with any component of the signal parallel to that. If implemented carelessly, this aggressive method will entail cosmological signal loss, although the problem is rectifiable.

Method 1 is described in Rao (1967), and is a simply a slight modification of the prescription in Sec. 3.2. Method 2 is the one we develop here and advocate for global 21 cm signal recovery.

We demonstrate the differences in approach using simple simulations. These have N_θ lines of sight and a number of




Figure 4. A simple simulation of 10 lines of sight with a random combination of three power laws.

power laws in the synchrotron emission (to limit the rank). The indices of these power laws are $\{-1.5, -1.8, -2\}$ and the amplitude is uniformly distributed and positive. Again, as amplitudes we take representative values from the recent observations of Voytek et al. (2014) and scale to the slightly higher frequencies of interest here. For the purposes here we will neglect thermal noise, and assume the integration is deep relative to the cosmological signal. Fig. 4 shows one realization of the foreground model. This model is only meant to be pedagogical, showing the main point of the algorithms here for independent sightlines in a simplified setting with three spectral modes. Beginning in Sec. 5 we will use the full extended GSM model data cubes described in Sec. 2.

3.4.1. Method 1: Foreground modes orthogonal to signal

To find foreground spectral modes that are orthogonal to the cosmological signal \mathbf{x} , we can find the largest eigenmodes of the restricted covariance

$$\hat{\Sigma}_{\perp \mathbf{x}} = [1 - \mathbf{x}(\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T] \hat{\Sigma}, \quad (21)$$

which in practice will be eigenmodes that represent foregrounds. By construction, any eigenmode \mathbf{f}_i of $\hat{\Sigma}_{\perp \mathbf{x}}$ will be orthogonal to \mathbf{x} , i.e. $\mathbf{f}_i^T \mathbf{x} = 0$. Let \mathbf{F}_{\perp} hold the restricted eigenvectors \mathbf{f}_i . The top left subplot of Fig. 5 shows the unrestricted eigenvectors of an input foreground simulation, while the bottom left subplot shows the restricted eigenvectors relative to a signal template.

Having identified our foreground eigenvectors, we can construct our estimator for the signal amplitude α using the methods introduced in Sec. 3.2. Once again, we first form a poor initial estimator of the signal, which we then adjust using the covariance between the poor estimator and the foregrounds (which is non-zero because a substantial foreground contamination remains in the estimate of the signal). The covariance of this estimator takes the form

$$\text{Cov}_{\hat{\alpha}, \hat{\beta}} = \begin{pmatrix} (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \hat{\Sigma} \mathbf{x} (\mathbf{x}^T \mathbf{x})^{-1} & (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T \hat{\Sigma} \mathbf{F}_{\perp} \\ \mathbf{F}_{\perp}^T \hat{\Sigma} \mathbf{x} (\mathbf{x}^T \mathbf{x})^{-1} & \mathbf{F}_{\perp}^T \hat{\Sigma} \mathbf{F}_{\perp} \end{pmatrix}, \quad (22)$$

which looks exactly like Eq. (12), but with all occurrences of \mathbf{Z} replaced by \mathbf{F}_{\perp} . This is unsurprising, since the modes in \mathbf{G} represent a diagonal piece of the covariance, $\sigma_{\text{inst}}^2 \mathbf{1}$, so are

not correlated in frequency with either the foreground modes or the signal.

Projecting out the piece of the signal estimator that is correlated with the foregrounds,

$$\hat{\alpha} = (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T (1 - \Pi) \bar{\mathbf{y}} \quad (23)$$

$$\Pi = \hat{\Sigma} \mathbf{F}_{\perp} (\mathbf{F}_{\perp}^T \hat{\Sigma} \mathbf{F}_{\perp})^{-1} \mathbf{F}_{\perp}^T. \quad (24)$$

That is, project foreground-contaminated modes in the data and then dot against the signal template. The operation Π puts the signal in the basis of foreground modes, weights those by their covariance, and moves back to the spectral space. Note that \mathbf{F}_{\perp} does not span the rest of the spectral space like \mathbf{Z} in the previous section. Here, Eq. (23) is not equivalent to $\hat{\alpha} = (\mathbf{x}^T \hat{\Sigma}^{-1} \mathbf{x})^{-1} \mathbf{x}^T \hat{\Sigma}^{-1} \bar{\mathbf{y}}$.

By construction, our estimator suffers from no loss of cosmological signal (despite the projecting-out of foregrounds) because $(1 - \Pi)\mathbf{x} = \mathbf{x} - \hat{\Sigma} \mathbf{F}_{\perp} (\mathbf{F}_{\perp}^T \hat{\Sigma} \mathbf{F}_{\perp})^{-1} \mathbf{F}_{\perp}^T \mathbf{x} = \mathbf{x}$. In the last step, we used the fact that $\mathbf{F}_{\perp}^T \mathbf{x} = 0$ through the choice of forming \mathbf{F}_{\perp} with eigenvectors restricted to be orthogonal to \mathbf{x} , Eq. (21).

The error on $\hat{\alpha}$ is

$$\text{Var}(\hat{\alpha}) \propto [\mathbf{C}_{\parallel, \parallel} - \mathbf{C}_{\parallel, \perp} \mathbf{C}_{\perp, \perp}^{-1} \mathbf{C}_{\perp, \parallel}] \quad (25)$$

$$\propto (\mathbf{x}^T \mathbf{x})^{-1} \mathbf{x}^T (1 - \Pi) \hat{\Sigma} \mathbf{x} (\mathbf{x}^T \mathbf{x})^{-1}. \quad (26)$$

Here $\mathbf{x}^T (1 - \Pi) \hat{\Sigma} \mathbf{x}$ can be interpreted as projecting all of the foreground modes orthogonal to the signal out of the (ν, ν') covariance and finding the noise with respect to that residual covariance. The residual covariance originates from foreground components parallel to the signal and from thermal noise—any (ν, ν') covariance component with a non-zero dot product into the signal \mathbf{x} will contribute to the variance of $\hat{\alpha}$. This property can be seen in Fig 5. The recovered signal has exactly the shape of the signal template \mathbf{x} , but the amplitude is dominated by foregrounds. This feature is clearly undesirable in the regime where foregrounds vastly exceed the signal.

Another undesirable property of this estimator is that the quoted error $\text{Var}(\hat{\alpha})$ relies on the foregrounds being normally distributed. Real foregrounds are strongly non-Gaussian, and we would like to avoid describing those components in the errors of $\hat{\alpha}$. We can make a simple modification to the estimator in Rao (1967) to treat both of these shortcomings.

3.4.2. Method 2: Signal component orthogonal to foregrounds

In some sense, the method described in Sec. 3.4.1 was too conservative. By limiting our labeling of foregrounds to modes that are orthogonal to the signal, we arrived at an estimator with no formal signal loss, but one in which substantial foreground residuals remained in the final answer. Here we consider an estimator that more aggressively projects out foregrounds. The result will be a lossy treatment, but we will also show how this can be rectified.

Instead of isolating the foregrounds orthogonal to the signal, consider the component of the signal that is orthogonal to the foregrounds. The signal is then partitioned as

$$\mathbf{x} = \mathbf{x}_{\parallel \mathbf{F}} + \mathbf{x}_{\perp \mathbf{F}} \quad (27)$$

For this decomposition to be meaningful, we must once again decide on a definition for our foreground modes. We let \mathbf{F} be the largest foreground eigenvectors of $\hat{\Sigma}$ (without the orthogonality restrictions of the previous section). In the calculation for $\hat{\Sigma}$, the mean $\bar{\mathbf{y}}$ contains the cosmological signal and is




Figure 5. Modes discovered from the (ν, ν') covariance of simple simulations with three synchrotron spectral indices (Fig. 4). These include only the foregrounds and simple signal model of Eq. (3) with $\Delta z = 0.5$ and $z_r = 10.6$. *Top left:* the largest three eigenvectors of the covariance $\hat{\Sigma}$ that fully describe the foregrounds. *Top right:* The filtered data $(1 - \mathbf{F}\mathbf{F}^T)\bar{\mathbf{y}}$. The input signal in red is largely recovered after cleaning the three foreground eigenmodes in the top left. In this toy model, the foregrounds are removed entirely, and only the signal (minus components parallel to the subtracted foregrounds) is left. *Bottom left:* The largest three eigenvectors of the restricted covariance $\hat{\Sigma}|_x = [1 - \mathbf{x}(\mathbf{x}^T\mathbf{x})^{-1}\mathbf{x}^T]\hat{\Sigma}$. These are all orthogonal to the cosmological signal. *Bottom right:* The data after subtracting the restricted eigenvectors [or $(1 - \mathbf{F}_\perp\mathbf{F}_\perp^T)\bar{\mathbf{y}}$] are exactly the signal template shape, by the construction of the mode functions. The amplitude, however, is clearly dominated by residual foregrounds. These are the foreground components parallel to the cosmological signal. While this procedure is immune from loss of cosmological signal, it is insufficiently aggressive when dealing with foregrounds, and the final result is strongly influenced by the foreground covariance. We note that the apparent recovery of the reionization step’s location in the bottom right subplot is spurious, for it is almost entirely driven by our choice of which theoretical template we use. The location of the “step” in the top right subplot, on the other hand, is truly constrained by the data, although signal loss means that it no longer appears as a clean step.

subtracted out. For this reason, the cosmological signal does not perturb the (ν, ν') covariance or its eigenmodes. This is in contrast to inhomogeneous 21 cm mapping, where the cosmological signal is stochastic, cannot be subtracted from the covariance, and does perturb the eigenvectors (Switzer et al. 2013; Masui et al. 2013). (The foreground modes can optionally be isolated from thermal noise by forming the (ν, ν') cross-variance of maps acquired at two different times, as was done in Switzer et al. 2013.)

To project the identified foreground modes out of the estimator, we apply $1 - \Pi = 1 - \mathbf{F}\mathbf{F}^T$. This operation splits the spectrum into a foreground contaminated and a (theoretically) foreground-clean subspace as

$$\begin{pmatrix} \Pi \\ 1 - \Pi \end{pmatrix} \mathbf{y} = \begin{pmatrix} \Pi \\ 1 - \Pi \end{pmatrix} \alpha \mathbf{x} + \begin{pmatrix} \mathbf{n}_{\text{fg}} \\ \mathbf{0} \end{pmatrix} + \mathbf{n}_{\text{inst}}. \quad (28)$$

Unlike the previous estimator Eq. (23), however, the foreground modes are no longer constructed to be orthogonal to the signal. By projecting out the foregrounds, then, it is likely that some cosmological signal will be lost. Mathematically, $(1 - \mathbf{F}\mathbf{F}^T)\mathbf{x}$ equals \mathbf{x}_\perp , not \mathbf{x} , and the estimator $\hat{\alpha} = (\mathbf{x}^T\mathbf{x})^{-1}\mathbf{x}^T(1 - \Pi)\bar{\mathbf{y}}$ is an incorrectly normalized estimator of the signal amplitude (i.e., it suffers from a multiplicative bias). Correcting this, we propose the revised estimator

$$\hat{\alpha} = \frac{\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\bar{\mathbf{y}}}{\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{x}}. \quad (29)$$

In the error calculation from Eq. (26), the key term $(1 - \Pi)\hat{\Sigma}$ now reduces to $1\sigma_{\text{inst}}^2$, under the assumption that all foregrounds are removed to a good approximation, and only Gaussian thermal noise remains. While the foregrounds may be strongly non-Gaussian, we will assume that the residuals after the modes \mathbf{F} have been subtracted are Gaussian and due to the thermal noise floor. This would need to be verified in analysis of real data.

The overall error is now

$$\text{Var}(\hat{\alpha}) \sim \left(1 + \frac{N_{\text{fg}}}{N_\theta - 1 - N_{\text{fg}}}\right) \frac{\sigma_{\text{inst}}^2}{[\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{x}]^2} \quad (30)$$

To get some rough intuition, assume that the foreground modes all have about the same overlap with the cosmological signal. Then $\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{x} \propto N_\nu - N_{\text{fg}}$. As N_{fg} approaches N_ν , the foreground modes do a progressively better job of spanning the N_ν spectral space. In the limit that they fully span the space (there is no signal distinguishable from foregrounds), the signal is also nulled and the error diverges. Alternately, as N_{fg} approaches N_θ , the foreground discovery process uses up all the spatial degrees of freedom and the error diverges. In this formulation, then, we are self-consistently including the possibility of signal-loss in our error analysis.

4. CONSTRAINTS ON PASSBAND CALIBRATION

The abstract design guidance suggested by this method is that 1) the experiment should not increase the rank of the foreground covariance (keep N_{fg} small) 2) if there are foregrounds with N_{fg} spectral degrees of freedom, at least $N_\theta = N_{\text{fg}}$ samples of its variations need to be observed if those foregrounds are to be subtracted.

A time-varying passband calibration is the worst instrumental systematic in this sense. Here each line of sight sees a slightly different foreground spectrum and so requires a different spectral function for cleaning. In the limit that each line of sight can have a different foreground, $N_{\text{fg}} \rightarrow N_\theta$ and the errors diverge. In contrast, a constant passband error only modulates the foreground covariance spectral functions, but it does not increase the rank of its covariance. Its impact is more aesthetic, can be corrected by re-calibrating off the brightest mode, and it does not fundamentally limit the investigation. (This is proven in Sec. 6 using specific notation developed there.)

Fig. 6 shows the 10 simulated lines of sight with a constant 10% calibration error. Here we multiply the spectrum along each line of sight by a constant $1 + \delta$, where δ is N_ν -long and normally-distributed. Fitting a smooth power law in this case would leave unacceptable residuals, but the modes are able to discover the constant calibration error and subtract the foreground. In contrast, Fig. 7 shows that even a 0.1% *variable* passband calibration error spreads foreground variance into many modes that are all significantly larger than the signal (10^{-5} there). To implement a variable passband, we multiply each line of sight by a different $1 + \delta_i$, representing a change in the instrument response with time.

The rule of thumb here is that if foregrounds are R times larger than the signal, then the per-pointing calibration needs to be stable to $\sim 1/R$. Averaging over sightlines tends to relax this constraint. In simulations shown in Sec. 7.1, we find that the stability must be better than $\sim 10^{-4}$ to have no effect. The constraint degrades rapidly above that. In the other limit of extremely high stability, we find that once calibration stability is suppressed below the thermal noise, there is no further gain in foreground discrimination.

An important point is that the method developed here can discover any response to foregrounds that varies across the sky (non-monopole). This does not necessarily need to originate from passband calibration or differential polarization response—these are just plausible sources of error from instrumental response.

5. INCORPORATING SPATIAL INFORMATION

The previous discussion assumed that the foregrounds are stationary and independent between different lines of sight. It is for this reason that the first step in the prescription outlined above was to form \bar{y} , the unweighted average of all lines of sight—with stationary statistics, there is no reason to prefer one direction of the sky over another, and an unweighted average provides the best signal-to-noise. This is the implicit assumption that is being made by experiments that average over large regions of sky at once (such as single-dipole experiments) and produce a single spectrum as their measurement. The results in preceding sections therefore also apply to experiments with no angular sensitivity, except without angular information, error bars and foreground cleaning methods cannot be informed by the data itself and must be derived *a priori*.

In reality, foregrounds are neither stationary nor independent. The synchrotron brightness varies across the sky (violating stationarity), and in addition the foregrounds are known to be spatially correlated (violating independence).

Angular information can be leveraged in several ways. First, non-stationarity of the foregrounds allows the estimator to down-weight parts of the sky that are known to be particularly noisy, such as the galactic center. Second, correlation information allows foreground properties in one part of the sky to be inferred from observations of another part of the sky. For example, in the unrealistically extreme limit of perfectly correlated foregrounds, a measurement of foregrounds in any part of the sky automatically allows a perfect prediction of foreground brightness in any other part of the sky.

To take advantage of angular information, L13 considered an optimal estimator given full-sky maps at all frequencies of interest, as well as a known $N_\nu N_{\text{pix}} \times N_\nu N_{\text{pix}}$ covariance \mathbf{N} between all pixels and all frequencies. Let \mathbf{d} be a vector of length $N_\nu N_\theta$ containing the measured sky maps. Here we use N_{pix} to denote the number of pixel indices in a full-sky healpix map (Górski et al. 2005). For reasons we will see in

Sec. 7.1, this also prevents confusion between the N_θ independent sightlines in the previous section and the present discussion, which is made more complex by beam convolution.

The observations \mathbf{d} are related to the global 21 cm spectrum \mathbf{s} via the measurement equation

$$\mathbf{d} = \mathbf{A}\mathbf{s} + \mathbf{n}, \quad (31)$$

where $\mathbf{n} \equiv \mathbf{n}_{\text{fg}} + \mathbf{n}_{\text{inst}}$ is the generalized noise, containing \mathbf{n}_{fg} and \mathbf{n}_{inst} as the foreground and instrumental noise contribution to \mathbf{d} , respectively. The $N_{\text{pix}} N_\nu \times N_\nu$ matrix \mathbf{A} is given by $\mathbf{1} \otimes \mathbf{e}_0$, where \otimes is the Kronecker product, $\mathbf{1}$ is an $N_\nu \times N_\nu$ identity matrix, and \mathbf{e}_0 is the spatial monopole, i.e., an N_{pix} -long vector of 1s. Its function is to copy, at every frequency, the value of the global spectrum \mathbf{s} to all the spatial pixels comprising the sky map. The maximum-likelihood estimate $\hat{\mathbf{s}}_{\text{ML}}$ for the full global spectrum is given by

$$\hat{\mathbf{s}}_{\text{ML}} = (\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1} \mathbf{A}^T \mathbf{N}^{-1} \mathbf{d}, \quad (32)$$

where $\mathbf{N} \equiv \langle \mathbf{n} \mathbf{n}^T \rangle - \langle \mathbf{n} \rangle \langle \mathbf{n} \rangle^T$ is the covariance matrix of the generalized noise, which is assumed to be known. To facilitate comparisons with previous work, we will begin by considering the estimator $\hat{\mathbf{s}}$ of the full global spectrum. Sec. 6 returns to the template amplitude constraint.

5.1. Spatial weighting to recover the spectrum

L13 used a series of analytic manipulations to show that incorporating angular information via Eq. (32) could in principle lead to large reductions in foreground contamination. In practice, however, this prescription is difficult to implement for a number of reasons. First, one may not possess sufficiently accurate models of the instrument and the foregrounds to write down the matrix \mathbf{N} , placing *a priori* expectations on the instrument response. This matrix cannot be derived from the data (unlike Σ from previous sections), since it has dimensions $N_\nu N_\theta \times N_\nu N_\theta$ and therefore contains many more degrees of freedom than the number of measurements $N_\nu N_\theta$. Moreover, even if the matrix is somehow known, its large size makes its inversion in Eq. (32) computationally challenging.

To deal with these challenges, consider a modified recipe where the sky maps are dealt with frequency-by-frequency [unlike in Eq. (32) where all frequencies and all lines of sight are mixed together] and an estimator for the global signal at each frequency is formed by spatially averaging with non-trivial weights [unlike in Eq. (6), where different lines-of-sight were equally weighted]. Let \mathbf{d}_β be a vector of length N_{pix} that represents the map within \mathbf{d} corresponding to the β th frequency channel. The global signal at this frequency channel can then be estimated by computing the weighted average

$$\hat{\mathbf{s}}_\beta = \mathbf{w}_\beta^T \mathbf{d}_\beta, \quad (33)$$

where \mathbf{w}_β is a vector of length N_{pix} with weight appropriate for the β th frequency channel. Throughout this paper, we adopt a convention where summations are written explicitly. The repeated indices therefore do *not* imply summations. The variance of this estimator—which we will seek to minimize—is

$$\text{Var}(\hat{\mathbf{s}}_\beta) \equiv \langle \hat{\mathbf{s}}_\beta^2 \rangle - \langle \hat{\mathbf{s}}_\beta \rangle^2 = \mathbf{w}_\beta^T \Phi_\beta \mathbf{w}_\beta, \quad (34)$$

where Φ_β is the $N_{\text{pix}} \times N_{\text{pix}}$ *spatial* covariance matrix of the map at the β th frequency. For our final estimator to be correctly normalized, we require the weights at each frequency




Figure 6. A toy model showing that an absolute passband calibration error that does not vary with time does not impact primordial signal recovery. *Top left:* Observations of the spectrum across 10 lines of sight with a stable 10% passband calibration error between bins (here neglecting thermal noise). The input synchrotron is taken from a mixture of three power laws, so is described by three spectral modes in the data. *Bottom left:* Foreground modes discovered in the data. Because of passband calibration error, these modes are not smooth. A constant calibration error does not increase the rank of the three input power laws. *Right:* Recovered cosmological signal compared to the input signal, and the cosmological signal with the three modes on the left removed. The fixed calibration error is cosmetic.




Figure 7. Residual foregrounds as a function of modes removed in a toy model with three foreground degrees of freedom. Amplitudes are normalized to 1 for zero modes removed, and in rms rather than variance. In the case with no variable calibration error, the variance in the spectrum quickly reaches thermal noise after three modes are removed. If the passband calibration varies between lines of sight, the residuals become significantly worse. Generally, for foregrounds 10^5 times larger than the signal, the passband should be stable to $\sim 10^{-5}$. Passband calibration instability spreads the variance of the three intrinsic foreground modes across many more degrees of freedom.

to sum to unity, i.e. $w_\beta^T e_0 = 1$. To obtain the minimum-variance solution for our weights, we use a Lagrange multiplier μ to impose our normalization constraint, and minimize the quantity

$$L = w_\beta^T \Phi_\beta w_\beta - \mu w_\beta^T e_0. \quad (35)$$

Doing so gives

$$w_\beta = \frac{\Phi_\beta^{-1} e_0}{e_0^T \Phi_\beta^{-1} e_0}, \quad (36)$$

and therefore

$$\hat{s}_\beta = \frac{e_0^T \Phi_\beta^{-1} d_\beta}{e_0^T \Phi_\beta^{-1} e_0}. \quad (37)$$

In words, these weights tell us to whiten the sky maps at each

frequency using the inverse spatial covariance before summing together different lines-of-sight and normalizing. Since the whitening takes the form of a non-diagonal matrix multiplication, this operation not only down-weights brighter parts of the sky, but also makes use of angular correlation information to better estimate the monopole signal. Tegmark et al. (2003) tackled the transpose of this problem for CMB foreground removal. Note that in modeling Φ_β , it is essential to capture the non-stationary nature of our galaxy's foreground emission. It is insufficient, for example, to describe the galaxy using an angular power spectrum alone, which assumes statistical isotropy. This would not only be physically unrealistic, but would also make our weights constant across the sky (a fact that can be derived by applying Parseval's theorem to our expression for w_β), defeating the purpose of using spatial weights in the first place.

5.2. Spatial weights within a separable covariance model

The spatial weight recipe that we have just specified can also be considered a special limit of the maximum-likelihood estimator, albeit with one small modification. Suppose that the full $N_\nu N_\theta \times N_\nu N_\theta$ covariance \mathbf{N} is separable, so that we can write it as $\Sigma \otimes \Phi$, where Σ is the $N_\nu \times N_\nu$ spectral covariance from previous sections, Φ is an $N_{\text{pix}} \times N_{\text{pix}}$ spatial covariance, and \otimes is the Kronecker product.

Using the identities $(\mathbf{G} \otimes \mathbf{H})^{-1} \equiv \mathbf{G}^{-1} \otimes \mathbf{H}^{-1}$ and $(\mathbf{G} \otimes \mathbf{H})(\mathbf{J} \otimes \mathbf{K}) \equiv \mathbf{GJ} \otimes \mathbf{HK}$, we have

$$\mathbf{A}^T \mathbf{N}^{-1} = \Sigma^{-1} \otimes e_0^T \Phi^{-1} \quad (38)$$

and

$$(\mathbf{A}^T \mathbf{N}^{-1} \mathbf{A})^{-1} = \Sigma \otimes (e_0^T \Phi^{-1} e_0)^{-1}. \quad (39)$$

Inserting these into Eq. (32), the final estimator in this separable approximation is then

$$\hat{s}_\beta \Big|_{\text{sep}} = \frac{e_0^T \Phi^{-1} d_\beta}{e_0^T \Phi^{-1} e_0}, \quad (40)$$

which is almost identical to Eq. (37). If we slightly relax the assumption of perfect separability by allowing Φ to acquire

a frequency dependence (so that $\Phi \rightarrow \Phi_\beta$), the correspondence becomes exact. Including this frequency-dependent non-separability also allows us to incorporate the fact that the instrumental beam will generally broaden toward lower frequencies, in addition to intrinsic non-separability of the foregrounds. Sec. 6.1 discusses non-separability caused by the differences between foregrounds and thermal noise.

Interestingly—but unsurprisingly—the spectral covariance Σ drops out of the optimal estimator for the global spectrum once separability is invoked. As discussed in L13, this is due to the fact that once we spatially average over the maps, we are left with a spectrum consisting of N_ν numbers. If our final goal is to measure a cosmological spectrum, that is also of length N_ν , the constraint that our estimator be correctly normalized and unbiased means that there is nothing left to do. Mathematically, this manifested itself in our derivation when the factors of Σ in Eqs. (38) and (39) canceled each other out when forming $(A^T N^{-1} A)^{-1} A^T N^{-1}$.

If we return to our method in previous sections, however, and attempt to constrain the amplitude α of a theoretical template, the spectral covariance Σ re-enters the discussion. When constraining α , one is seeking to measure a single number from N_ν measurements, which in general ought to be combined non-uniformly if these measurements have different error properties (as captured by Σ).

6. ANALYZING GLOBAL 21 CM DATA CUBES

In this section we combine the finite-rank, *in situ* estimate of spectral foregrounds with a model of angular correlations. The end product will be an estimator for the amplitude of a cosmological signal template that can be applied to experimental global 21 cm data cubes. The angular part of the estimator uses assumed spatial distributions to improve the amplitude estimate. In contrast, the central aspect of this estimator is that it does not assume a set of *spectral* modes, or spectral covariance of the data.

It is notationally convenient to write the $N_{\text{pix}} N_\nu$ -long data matrix \mathbf{d} by reshaping it into an $N_\nu \times N_{\text{pix}}$ matrix \mathbf{Y} . This “data cube” is a stack of maps at all the observed frequencies. The original data vector \mathbf{d} can be recovered through the serializing “vec()” operation $\text{vec}(\mathbf{Y}) = \mathbf{d}$. This data cube form naturally accommodates weighing operations by the separable covariance. A left-multiplication acts on the spectral direction, and a right-multiplication acts on the spatial direction.

The essential weighting operation on the full dataset is $N^{-1} \mathbf{d}$, which under the separable covariance assumption becomes

$$N^{-1} \mathbf{d} = (\Sigma \otimes \Phi)^{-1} \mathbf{y} = \Sigma^{-1} \mathbf{Y} \Phi^{-1} \quad (41)$$

using identities of the Kronecker and vec operations.

A model for the noise on \mathbf{Y} is to draw some matrix \mathbf{E} with shape $N_\nu \times N_{\text{pix}}$ of normal deviates from $N(\mu = 0, \sigma^2 = 1)$, and then correlate the spectral and spatial parts as $\Sigma^{1/2} \mathbf{E} \Phi^{1/2}$ to give a noise realization drawn from the covariance $\Sigma \otimes \Phi$.

The data model in this presentation is

$$\mathbf{Y} = \alpha \mathbf{x} \mathbf{e}_0^T + \Sigma^{1/2} \mathbf{E} \Phi^{1/2}. \quad (42)$$

This is a specialized case of the “growth curve” model (Kollo & von Rosen 2005), reviewed in Appendix B. We can transfer the separable method developed in the Sec. 5.2 to the present case of constraining the template amplitude α by letting $s \rightarrow \alpha$ and $\mathbf{A} = \mathbf{1} \otimes \mathbf{e}_0 \rightarrow \mathbf{x} \otimes \mathbf{e}_0$ (since we are

now assuming a form for \mathbf{x} as a theoretical template). Here, the maximum-likelihood estimate for the amplitude of some spectral template \mathbf{x} is

$$\hat{\alpha} = (\mathbf{x}^T \Sigma^{-1} \mathbf{x})^{-1} \mathbf{x}^T \Sigma^{-1} \mathbf{Y} \Phi^{-1} \mathbf{e}_0 (\mathbf{e}_0^T \Phi^{-1} \mathbf{e}_0)^{-1}. \quad (43)$$

This applies the familiar maximum likelihood estimators for the spatial monopole (with respect to covariance Φ) and spectral template (with respect to covariance Σ) to the right and left side of our data cube \mathbf{Y} .

Following Sec. 3.2, we can seek to replace Σ with a (ν, ν') covariance measured in the data, $\hat{\Sigma}$. Kollo & von Rosen (2005) show that, again, this choice of weight yields the maximum likelihood $\hat{\alpha}$. Again note that the errors on $\hat{\alpha}$ will be inflated (and non-Gaussian) by the fact that the estimation uses additional degrees of freedom in the data. As before, we want to find the sample variance with the monopole mean spectrum subtracted so that the cosmological signal does not enter the covariance. Additionally, now that the angular covariance Φ^{-1} is known, it can be applied in the covariance estimation as

$$\begin{aligned} \hat{\Sigma} &= \mathbf{Y} \Phi^{-1} [\mathbf{1} - \mathbf{e}_0 \mathbf{w}^T] \mathbf{Y}^T \\ \mathbf{w} &= \Phi^{-1} \mathbf{e}_0 (\mathbf{e}_0^T \Phi^{-1} \mathbf{e}_0)^{-1} \end{aligned} \quad (44)$$

where \mathbf{w} is an N_{pix} -long weight map.

The estimator for the monopole spectrum template amplitude is then

$$\hat{\alpha} = (\mathbf{x}^T \hat{\Sigma}^{-1} \mathbf{x})^{-1} \mathbf{x}^T \hat{\Sigma}^{-1} \mathbf{Y} \mathbf{w} \quad (45)$$

The inverse Φ^{-1} of the $N_{\text{pix}} \times N_{\text{pix}}$ spatial covariance appears explicitly only in the construction of the sample variance. Elsewhere, it is collapsed onto the single simple weight map \mathbf{w} . One approximation that can be made in practice is to replace Φ^{-1} with its diagonal, which simply weighs angles differently in the (ν, ν') variance estimation.

Similar to the discussion in Sec. 3.3, we can replace the $\hat{\Sigma}^{-1}$ operation by a projection that removes a finite number of contaminated spectral modes. This is equivalent to 1) identifying a data subspace parallel to foregrounds that is contaminated and fully removed, 2) keeping a data subspace orthogonal to foregrounds.

Let \mathbf{F} contain the largest eigenvectors of the sample covariance $\hat{\Sigma}$. Then

$$\hat{\alpha} = \frac{\mathbf{x}^T (\mathbf{1} - \mathbf{F} \mathbf{F}^T) \mathbf{Y} \mathbf{w}}{\mathbf{x}^T (\mathbf{1} - \mathbf{F} \mathbf{F}^T) \mathbf{x}}. \quad (46)$$

The denominator $\mathbf{x}^T (\mathbf{1} - \mathbf{F} \mathbf{F}^T) \mathbf{x}$ is a scalar normalization.

This is analogous to Eq. (29), except that instead of a simple $\bar{\mathbf{y}}$, each spatial slice of the map \mathbf{Y} is weighted to find the mean as $\mathbf{Y} \mathbf{w}$.

Using this notation, we can easily prove that constant passband calibration errors do not increase the rank of foreground spectral modes. Let $\mathbf{1} + \delta \mathbf{c}$ be an N_ν -long vector that represents constant mis-calibration of the passband. If a constant calibration error multiplies all lines of sight, $\mathbf{C} = \text{diag}(\mathbf{1} + \delta \mathbf{c})$ modifies the data as $\mathbf{C} \mathbf{Y}$. So long as the response is not zeroed out at some frequencies, \mathbf{C} is invertible, and $\text{rank}(\mathbf{C} \hat{\Sigma} \mathbf{C}^T) = \text{rank}(\hat{\Sigma})$ for invertible \mathbf{C} . Intuitively, a constant passband error can be thought of as just a per-frequency rescaling of units. Since such a rescaling does not affect the rank of the foregrounds, there is no increase in the number of foreground modes that need to be constrained.

Our self-discovery scheme for spectral modes will therefore be just as effective when constant passband mis-calibrations are present.

6.1. Deriving the final estimator: a two-component covariance

In deriving Eq. (46), the key assumption that we made was that the total covariance is separable into spatial and spectral components. This confuses foregrounds and thermal noise, which each have very different angular distributions. For example, we know that a synchrotron spectral mode is associated with broad spatial distributions of galactic emission. In contrast, thermal noise is always uncorrelated across angular pixels, but may change in variance across the sky. In the limit that spectral foreground subtraction is perfect, only Gaussian thermal noise will remain. In contrast, the separable assumption would dictate that it has approximately the same angular correlations as the galaxy. This issue ultimately translates into the angular weight w that should be applied.

The solution to this problem is to let the total covariance be non-separable, but comprised of independently separable foreground and noise components. Here, $N = N_f + N_g$ where $N_f = \Sigma_f \otimes \Phi_f$ and $N_g = \Sigma_g \otimes \Phi_g$. For concreteness, imagine that the “f” component consists of modeled foreground modes in F , while the “g” component consists of any residuals after the projection operation. This is exactly analogous to the F/G split in Rao (1967) described in Sec. 3.3. If our discovered foreground modes happen to perfectly capture the true foregrounds, then the “g” component would consist only of Gaussian thermal noise. In general, however, we need not make this assumption.

The central problem with the model $N = \Sigma \otimes \Phi$ is that $(1 - FF^T)$ only acts on the Σ part, but not Φ . In contrast, in the two-component model, when $(1 - FF^T)$ acts on the data, the remaining covariance is only $\Sigma_g \otimes \Phi_g$, e.g., it correctly ascribes Φ_f ’s spatial correlations to the modes that F removes.

We can now reassess the w in Eq. (46) which yields the minimum-variance estimate of α . To clarify the discussion, let $q \equiv [x^T(1 - FF^T)x]^{-1}(1 - FF^T)x$. The first step of Eq. (46) in this notation is to form $q^T Y$, which can be interpreted as a series of estimators for α , one for each line-of-sight.

With our two-component covariance, a little algebra reveals that the variance to minimize is given by

$$\text{Var}(\hat{\alpha}) = (q^T \Sigma_f q)(w^T \Phi_f w) + (q^T \Sigma_g q)(w^T \Phi_g w). \quad (47)$$

The first term is zero by construction, since $q^T \Sigma_f \propto (1 - FF^T)F = 0$. Minimizing the remaining term subject to the constraint that the weights sum to unity requires minimizing the quantity

$$L = (q^T \Sigma_g q)(w^T \Phi_g w) - \mu w^T e_0, \quad (48)$$

where μ is again a Lagrange multiplier. Since $q^T \Sigma_g q$ is just a constant, it can be absorbed into μ by a simple rescaling. One then sees that L is identical to Eq. (35) except that Φ_g (the spatial covariances of the residuals) takes the place of Φ_β (the full spatial covariance at the β th frequency). We may thus import our previous solution for the spatial weights to find

$$w = \Phi_g^{-1} e_0 (e_0^T \Phi_g^{-1} e_0)^{-1} \quad (49)$$

In conclusion, we see that a simple re-interpretation of Eq. (46) is all that is required to accommodate a two-component

covariance model. The new weights w reflect the spatial covariance of residuals after the frequency mode subtraction rather than the spatial covariance of the sky itself. In spirit, this is reminiscent of the approach in L13, where a best-guess model of the foregrounds was first subtracted from the sky maps before the variance of the residuals was minimized. The difference here is that our best-guess cleaning is based on the data rather than a model. In addition, the estimator operates with separable spectral and spatial steps that are computationally trivial (a feature that we have managed to retain even with the *non*-separable two-component covariance of this section).

6.2. Angular weighting

If the F modes cleanly separate the finite rank foregrounds, the spatial covariance Φ_g that determines the angular weight represents Gaussian thermal noise. In this case Φ_g is diagonal, and proportional to T_{sys}^2 based on the radiometer equation. Experimentally, synchrotron emission on the sky dominates the receiver noise temperature, so the angular weight w is simply to divide by the synchrotron template-squared.

In reality, there will not be a clean separation of foregrounds and thermal noise. This is related to the fact that there is not a rigorous way to determine the number of modes to remove (Sec. 7.4). If the eigenvalue spectrum has a long tail, we can remove the obvious, high signal-to-noise modes, but some residual will remain. This means that the w needs to represent the angular covariance of the residual foregrounds. Following L13, we split the problem of non-stationarity and angular correlation of the angular covariance by writing

$$\Phi_g = \text{diag}(m) Q \text{diag}(m)^T \quad (50)$$

where m is the foreground mean map. This is equivalent to dividing the foreground by the expected mean, and then encoding the angular correlations of that normalized map in the matrix Q . Let these correlations be full-sky, isotropic and diagonal. If T is a matrix that converts from real space to spherical harmonics, then the correlations will be diagonal $\Lambda = TQT^T$, where Λ is just some function of ℓ . The operation Φ^{-1} can be interpreted algorithmically as

- Divide each spatial slice of the data by a foreground model map (a synchrotron template)
- Transform to spherical harmonic space and de-weight ℓ -by- ℓ by w_ℓ (defined below)
- Transform back to map space and divide by the foreground model map again.

Here the ℓ weighting reflects the angular correlations of any residuals after the $(1 - FF^T)$ cleaning step. This is also the step where galactic masks could be applied in real space through a choice of m^{-1} that is zero in those regions.

The choice of w is related to the optimality of the estimate of $\hat{\alpha}$. As a first pass, we can always use the Gaussian thermal noise assumption, which has $w_\ell = 1$. An ℓ -by- ℓ weighting of the residuals then has the potential to improve the estimate of $\hat{\alpha}$. Detecting some angular power spectrum in the $(1 - FF^T)$ -cleaned map that is measurably different than Gaussian noise could indicate that a better weight w_ℓ is possible. A simple, two-parameter model for the correlations is to let the residuals have some amplitude ξ relative to thermal noise and correlation length σ , as

$$w_\ell = \left[\xi e^{-\sigma^2 \ell(\ell+1)/2} + e^{\theta_b \ell(\ell+1)} \right]^{-1}. \quad (51)$$




Figure 8. *Top:* Spatial weights at 80 MHz for noise-dominated residuals ($\xi = 0.5$; $\sigma = 20^\circ$). The main purpose of spatial weighting in this regime is to average down noise, so all the weights are positive. Bright regions of the sky (which cause more instrumental noise for sky-noise dominated instruments) are down-weighted. *Bottom:* Spatial weights at 80 MHz for foreground-dominated residuals ($\xi = 50$; $\sigma = 20^\circ$). In this regime, the spatial weights play a role in foreground subtraction, and thus the weights are both positive and negative, as different parts of the sky are differenced to mitigate residual foregrounds. The beam size assumed in both cases is $\theta_b = 10^\circ$. The weights are defined to sum to 1, but only the shape, not normalization, is relevant.

Here we have assumed that the input data are deconvolved by the beam before applying angular weights (dividing by B_ℓ). The term $e^{\theta_b \ell(\ell+1)}$ describes thermal noise in that case. By picking the σ and ξ to give the lowest error, we may give preference to parameters where the error spuriously scatters low. A better alternative is use the known instrumental thermal noise to inform ξ , and assume residuals in the data are correlated across the beam scale θ_b (not determining w_ℓ based on the error of $\hat{\alpha}$). Ideally, the residuals would be dominated by nearly Gaussian thermal noise, and $w_\ell = 1$ would be a sufficiently optimal weight.

In Fig. 8 we show example spatial weights computed from Eqs. (49), (50), and (51) in two different regimes. The top panel shows a noise-dominated case ($\xi = 0.5$) while the bottom panel shows a foreground residual-dominated case ($\xi = 50$). In the noise-dominated case the spatial weights are designed to simply average down the noise, and are therefore all positive. The galactic plane and the bright point sources are seen to receive almost zero weight, since they contribute the most thermal noise in a sky-noise dominated instrument. In the foreground-dominated case, the spatial weights attempt to further suppress residual foreground contamination. There are thus both negative and positive weights, since different parts of the sky can be differenced to subtract off foregrounds. In both regimes, the right half of the sky is given more weight, since the galaxy is dimmer there (see Fig. 2).

An important property of our spatial weighting scheme is that the normalization is irrelevant, and only the shape mat-

ters. This is a particularly attractive property, since most models of the sky at our frequencies of interest are based on interpolations and extrapolations from other frequencies. The amplitude may have large uncertainties, but the shape information is likely to be much more reliable. If one is particularly confident about the available shape information, it may be prudent to go one step further and to set $w_{\ell=0} = 0$ by hand. Examining Eqs. (50) and (51) reveals that this projects out any component of the measured sky that has precisely the same spatial shape as our foreground model. Note that even though this is the $\ell = 0$ weight, we do not destroy the global monopole signal that we seek to measure, since the w_ℓ act in a pre-whitened space *after* dividing by m . The global signal therefore takes the form $1/m$ and is no longer just the $\ell = 0$ mode. It is instead spread out over a wide range of ℓ values. Setting $w_{\ell=0} = 0$ then sacrifices the sensitivity in one of the modes within this range, resulting in slightly increased error bars as one deviates slightly from the optimal prescription described above. However, this may be a cost that is worth bearing for the sake of robustness in an aggressive campaign against foreground systematics.

6.3. Summary of the proposed algorithm

The final algorithm suggested here for analyzing global 21 cm data cubes is

- Use a spatially-weighted average to project the monopole spectrum out of the map and find the (ν, ν') sample covariance.
- Find the largest eigenvectors of the sample covariance and project them out of the map.
- Combine lines of sight using a prescription for the angular correlation of the residuals (which conservatively is to divide by the synchrotron template-squared).
- Dot this against the signal template to find the amplitude and perform error analysis.

While this algorithm is intuitive, our goal here has been to describe some of the implicit choices: 1) discovering frequency modes within the data (and implications for errors), 2) choosing to remove the part of the signal parallel to the foreground modes, 3) assuming a two-component separable form for the thermal noise and foregrounds/respectively. The next section describes several considerations for using estimators of this type, and challenges in global 21 cm signal measurement.

7. CONSIDERATIONS FOR ANALYZING GLOBAL 21 MAPS

In the next few subsections, we will consider 1) the notion of resolution elements and the choice of beam size, 2) combined passband stability and resolution considerations, 3) challenges in error estimation from residual monopoles, 4) determination of the number of modes to remove, 5) extensions to simple amplitude constraints, 6) applications of our methods to the pre-reionization absorption feature, and 7) complications in foreground mitigation that result from Faraday rotation.

7.1. Resolution elements and foreground modes

In Sec. 3, the data were spectra of independent lines of sight. There, the number of foreground modes removed (N_{fg})

compared to number of sightlines (N_θ) entered the errors as $\sim (N_{\text{fg}} - N_\theta)^{-1}$. The sense of N_θ is more complex in our present case of an $N_\nu \times N_{\text{pix}}$ \mathbf{Y} data cube where the spatial part of sky signal is convolved by an instrumental beam.

Intuition suggests that the number of independent sight lines is approximately the number of beam spots on the sky, but this is incorrect. If each spatial slice of a noiseless full-sky map is written in terms of its spherical harmonics $T_{\ell m}$, convolution by the beam is multiplication by B_ℓ . Unless the beam falls to zero somewhere, this operation is invertible and so does not modify the rank of $\hat{\Sigma}$. The rank of \mathbf{Y} is N_{fg} so long as $N_{\text{pix}} > N_{\text{fg}}$ and $N_\nu > N_{\text{fg}}$. Hence, all the spectral modes can still be recovered from the noiseless beam-convolved map.

With wider beams, higher eigenvalue spectral eigenmodes (typically the less smooth modes) are washed out by the beam. These modes do not disappear entirely, but their amplitude diminishes. We can develop a rough rule for this effect by finding the eigenvalue spectrum of full-rank white noise that is convolved by a given beam size. A fitting form for the suppression of eigenvalues (normalized to give no suppression for infinitely fine beams) is $b_n \approx 10^{-(\theta_{\text{FWHM}}/300^\circ)^2 f_{\text{sky}}^{-1} n}$, as a function of eigenvalue n and full-width at half-max (FWHM) of the beam.

At first sight, it would appear that a more rapidly falling eigenvalue spectrum is beneficial because it means that the foregrounds are better described by fewer covariance modes. When the suppression occurs because of beam convolution, however, it can be problematic in the presence of instrumental noise. With instrumental noise, the information about a mode may be suppressed below the noise floor. That foreground can therefore no longer be cleanly detected and removed, and may contaminate the signal and lead to larger errors. Such detection and removal is potentially crucial, since higher eigenvalues can still greatly exceed the signal (even after beam suppression) because of the hierarchy of scales.

We therefore come to the conclusion that even if angular resolution does not impose a hard limit on the number of foreground modes that may be discovered (as suggested in the case of independent sight lines), it does put a practical limit on the efficacy of the cleaning. Note that the eigenvalues of $\Sigma = \Sigma_{\text{fg}} + \Sigma_{\text{inst}}$ will be different than Σ_{fg} . Adding noise has the effect of mixing the high-variance foreground modes with the thermal noise modes, making it harder to cleanly separate foregrounds.

A rule of thumb for finding the number of modes discoverable before they are confused by noise is to plot the eigenvalue spectrum of noiseless foregrounds and the thermal noise, then find where they intersect. Modes some margin above this point will still be well-measured. Consider a normalized eigenvalue spectrum of foregrounds that follows $\lambda_n \approx b_n \cdot 10^{-an}$ (Liu & Tegmark 2012). The eigenvalues of thermal noise are relatively flat in n in comparison. Finding the intersection with the noise level, the number of foreground modes that can be learned in the data is

$$N_{\text{fg}} \sim \log(\text{FNR}) \left[\left(\frac{\theta_{\text{FWHM}}}{300^\circ} \right)^2 f_{\text{sky}}^{-1} + a \right]^{-1}, \quad (52)$$

where FNR is the foreground-to-noise ratio (in map variance). In our simple models here, a may be a factor of a few, for only a few modes in total. In this idealized setting, the rather lenient scale of 300° in our fitting form for b_n means that an




Figure 9. Error of the recovered cosmological template amplitude as a function of calibration stability and instrumental resolution, where the fiducial amplitude is $\alpha = 1$. The thermal noise is set so that in the limit of no calibration or beam effects, the amplitude can be constrained to 5σ . Here we consider foregrounds with four intrinsic spectral degrees of freedom. The constraint degrades for FWHM above $\sim 50^\circ$ and 10^{-4} stability in passband calibration. At low resolution, the modes can not be discovered. Calibration stability worse than the thermal noise produces many more modes that must be constrained and removed. The Monte Carlo estimates in the upper right become very noisy and we saturate the color scale.

instrument could cover the full sky at poor resolution and still discover all the foreground modes. In reality, instrumental effects will drive a higher number of required modes. For example, Switzer et al. (2013) had to remove tens of modes to clean foregrounds even though only approximately four intrinsic synchrotron modes were expected.

An advantageous factor here is that the sky signal is convolved by the beam while instrumental thermal noise is not. Hence, thermal noise can be estimated within the data using N_{pix} independent samples. In an ideal experiment, all of the foreground modes can be discovered and subtracted in \mathbf{F} so that only thermal noise remains. Then errors could then be assessed from the map's thermal noise, and signal loss could be corrected against the subtracted modes as $\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{x}$.

It is generally beneficial to cover larger areas with finer resolution to better constrain and discover foreground modes. Higher resolution also helps focus de-weighting of particular contaminated sky regions. An important exception to this occurs when the number of foreground modes scales with the FWHM or sky fraction. Both Faraday rotation and passband instability will generally increase the number of modes needed for studies on larger areas with smaller beams. Each new beam that is observed could see a new rotation measure, or be observed with a different passband.

7.2. Summarizing the effects of passband stability and angular resolution

Fig. 9 summarizes two of the main points of our paper. Here, we use the galaxy model described in Sec. 2 to perform a Monte Carlo simulation for the errors on α as a function of calibration stability and instrumental resolution. Fig. 10 shows the four largest modes of this extended GSM model. All fur-




Figure 10. Spectral eigenmodes of the foreground model described in Sec. 2. The second, third, and fourth modes are 3×10^{-4} , 8×10^{-7} and 6×10^{-8} down from the first mode (in variance), respectively.

ther modes of the model itself (without noise) are negligible, at the level of machine precision. Note that in contrast to the simple pedagogical models in Sec. 3 (with three modes), the full data cube simulation is more representative of the synchrotron sky.

The sense of the passband calibration stability in our Monte Carlo is pixel-to-pixel (e.g., not beam convolved), so is not affected by the beam convolution operation. We set the thermal noise level to produce a 5σ detection of α in the absence of beam and calibration effects. For beams that are too large, not all foreground modes can be recovered and subtracted. In the simulations here with four modes, the constraint on α worsens gradually as the FWHM exceeds $\sim 50^\circ$. For calibration stabilities worse than $\sim 10^{-4}$, each line of sight responds differently to bright foregrounds, so requires a new spectral function to remove. As soon as these passband calibration perturbations exceed the noise floor, they rapidly degrade the constraint.

Decreasing the thermal noise through longer integration times loosens the beam resolution requirement (because foreground modes are better measured). It does not modify the passband stability requirement. Passband stability is driven by the separation in scales between the foregrounds and cosmological signal, and the fact that each line of sight has a different response to the foregrounds. Even if the modes are well-measured relative to thermal noise, they quickly overwhelm the signal.

Throughout, we have assumed that the angular resolution at all frequencies is constant. This is difficult to arrange experimentally. More typically, $\text{FWHM} \propto \lambda$ by diffraction. In the frequency range considered here, a diffraction-limited FWHM would change by a factor of 2.5. Variation of the beam will mix spatial structure into frequency structure as different beams at each frequency see different angular contamination patterns (Morales et al. 2012). Implicitly, we have assumed that the beam is known well enough to convolve all the frequency slices to a common resolution. In contrast, if no correction is performed, the simulations described in Sec. 2 with four modes will proliferate into 11 modes. These new spectral modes describe the spatial structure that is mixed in to the spectral direction. The frequency-dependent beam must therefore be corrected to achieve efficient cleaning. Errors in the beam model executing this correction then translate into new spectral contamination modes.

The tolerance of beam models is experiment-dependent and is beyond the scope of the present work, but must be carefully considered. We emphasize that frequency-dependent beam effects impact any global 21 cm experiment, whether

they choose to take advantage of angular information or not. The new spectral mixing modes reflect a best effort of our algorithm to describe new spectral covariance that must be cleaned to reach the 21 cm monopole. Careful experimental design and measurement of the beam will control these covariance modes (which remain inaccessible in a single-beam measurement).

Here we have limited the intrinsic foreground rank to four. In reality, it may be discovered that foregrounds on the sky have higher rank, pushing the resolution constraint. Also, we assume a full sky here—an experiment covering a smaller fraction of the sky would need f_{sky}^{-1} more resolution in our simple approximations here.

7.3. Challenges in error estimation

In the absence of prior information, monopole contamination is indistinguishable from the signal. This contamination may arise from a synchrotron monopole, or from additive $T_{\text{rx}}(\nu)$ of the receiver that is temporally constant but has some spectral structure. The foreground modes are estimated from the (ν, ν') sample covariance $\hat{\Sigma}$ of the map with the mean spectrum removed. The modes in \mathbf{F} represent the frequency components of the spatially fluctuating part of the foregrounds, and the operation $1 - \mathbf{F}\mathbf{F}^T$ is uninformed by a monopole signal.

The monopole foregrounds may have significant overlap with the spectral functions of the spatially fluctuating foregrounds. In our simulations here, the foregrounds are randomly drawn from a fixed set of spectral modes. For example, if there is spatially-varying synchrotron with a spectral index of -2 , that synchrotron emission will also be subtracted from the monopole by $1 - \mathbf{F}\mathbf{F}^T$.

In reality, we cannot rule out the possibility of a foreground monopole that remains even after cleaning the spatially-varying synchrotron modes. Further, the spectral pattern of constant instrument $T_{\text{rx}}(\nu)$ will never be discoverable by mapping across the sky, and could be confused with the global signal. These present serious challenges for rigorous error estimation.

Performing the angular operations on the right of Eq. (46) first, the noise terms of $\hat{\alpha}$ can be isolated as

$$\hat{\alpha} = \alpha|_{\text{true}} + \frac{\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)}{\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{x}}(\mathbf{n}_{\text{fg}} + \mathbf{n}_{\text{inst}}) \quad (53)$$

where \mathbf{n}_{fg} and \mathbf{n}_{inst} are the foregrounds and thermal noise of the monopole through $\mathbf{Y}\mathbf{w}$. We lump the constant $T_{\text{rx}}(\nu)$ in with \mathbf{n}_{fg} rather than introducing a new term.

While the thermal noise contribution $\propto \mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{n}_{\text{inst}}$ has a well-defined distribution and can easily be included in the errors, $\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{n}_{\text{fg}}$ is much more challenging. If $\mathbf{x}^T(1 - \mathbf{F}\mathbf{F}^T)\mathbf{n}_{\text{fg}}$ were always positive, then $\hat{\alpha}$ could be interpreted as an upper bound. This practice is common in treating foregrounds to the power spectrum because their contribution is always positive in the quadratic quantity (Switzer et al. 2013; Parsons et al. 2013; Dillon et al. 2014; Paciga et al. 2013). Here, we have no guarantee that the residual foreground monopole dotted into the signal is positive. Recall that even though synchrotron may be a strong power law, the residuals here will largely represent unknown instrumental factors that could dot with arbitrary sign with the signal template. It has long been known that higher-order polynomials lead to better subtractions. This suggests that resid-

uials will generally produce two-sided errors. Further, any monopole residual foregrounds are likely to be non-Gaussian, inheriting from non-Gaussianity of the bright foregrounds.

These “monopole” foregrounds do not strictly need to be a monopole on the sky—they only have to be constant over the area of the survey, if that only covers part of the sky. This again bolsters the argument that global 21 cm experiments should observe as much of the sky as possible with modest resolution to maximize their sensitivity to spatial variations that could lead to significantly cleaner signal recovery.

Barring any prior information about the foregrounds or instrumental response, the best ways to immunize the analysis against this response are 1) to document that the spectral cube cleaned with $(1 - \mathbf{F}\mathbf{F}^T)$ is consistent with Gaussian thermal noise, 2) to cover a large area of the sky with modest resolution to capture additional spatial variation of foregrounds. Additional information about spectral smoothness and instrumental response can be used (Bowman & Rogers 2010; Voytek et al. 2014) to further clean the spectral monopole. Quoted errors would be based only on thermal noise, and boosted to account for signal lost in the cleaning process, Eq. (30). Another possibility would be to conduct e.g. surveys of the North and South Celestial Poles with different receiver architectures. If these were performed independently and had the same result up to thermal errors, it could lend some confidence that the any residual monopoles in the respective surveys are small compared to noise.

7.4. Determining the number of modes to remove

Fig. 11 shows the fractional rms of the cosmological reionization signal remaining after removing the four foreground modes of the extended GSM (Sec. 2), shown in Fig. 10. We use the operation $(1 - \mathbf{F}\mathbf{F}^T)$ described previously. The foreground spectral functions overlap strongly with the signal, especially for slower, spectrally smoother reionization histories (high Δz in the plot). Nulling four foregrounds causes significant signal loss over a range of reionization models. Any instrumental response that increases the rank of these foregrounds will result in additional signal loss from the new modes it introduces.

In the case of these simulations, we know the precise number of spectral degrees of freedom. In experimental practice, the number of contaminated modes can not be determined directly. This can be made especially ambiguous by an eigenvalue spectrum that drops slowly rather than reaching a clear noise floor.

A central challenge of any global 21 cm signal estimator is the determination of the number of modes to remove (Rao 1967; Kollo & von Rosen 2005). In similar methods applied to the 21 cm autopower (Switzer et al. 2013), an advantage is that the bias from foregrounds is purely additive. In this case, an experiment can report the constraint with errors as a function of modes removed. This will generically fall as more contaminated modes are removed. This decline will level out if the foreground have been discovered. As further modes are removed, the errors will increase due to cosmological signal loss, which is self-consistently included in our formalism. A reasonable prescription for a bound is to find the number of modes which gives the most stringent upper bound.

In the global spectrum case, foreground residuals after N spectral modes are removed can have either sign in the dot product with the signal template. This two-sided bias should stabilize as most of the foreground modes are discovered. For




Figure 11. Fraction of the root-mean-squared of the reionization signal remaining after the first four foreground modes are removed. The signal is given by Eq. (3) for a simple reionization scenario with central redshift z_r and width Δz . The first mode is a power law (see Fig. 10) and most of the signal remains for these reionization scenarios. The other three modes carry more redshift information, and tend to express signal variations at high redshift better. In all cases, as the signal becomes smoother (Δz increasing) it becomes harder to differentiate from the foreground modes.

more modes removed, the errors will grow, but the central value should remain stable if most of the foreground covariance is described and removed. The summary of a global 21 cm experiment results should include plots of the spectral modes removed and characterization of the constraint as a function of modes removed.

Both this issue and the residual monopole in Sec. 7.3 cannot be addressed rigorously. An irreducible challenge of global 21 cm measurement is determining whether any foreground bias remains. The methods here provide excellent prospects for cleaning the data and guiding an experiment, but do not solve these central issues.

7.5. Extending the spectral template

We have discussed constraints on the amplitude of a cosmological 21 cm signal template that is known in advance. The primary goal of the current generation of instrument is discovery of this amplitude. Subsequent experiments will then constrain the spectrum itself. The methods developed here can also be applied to that case.

There is a continuum of models between the template estimate here and a full spectrum. One possibility would be to use prior information about the redshift of reionization and estimate two amplitudes, $\alpha(z < z_r)$ and $\alpha(z > z_r)$ then examine the significance of a step feature. Another option would be to constrain a fixed number of spectral spline amplitudes. In the limit that each frequency bin has a different amplitude, the methods developed here can find the component of that signal orthogonal to the identified foreground modes.




Figure 12. The same as Fig. 12, except that we now consider a pre-reionization absorption feature parameterized by a Gaussian. The duration is given as the σ -width of the Gaussian. Note that significantly more signal remains.

The decision between models must juggle 1) use of prior information, 2) discriminatory power, 3) insight about the trajectory of the global 21 cm signal. This model selection problem is deferred to future work.

7.6. Application to pre-reionization absorption

Throughout the paper, we have considered the global 21 cm signal from reionization. This choice was primarily pedagogical, to have a simple step around which the foregrounds are subtracted. The maximum contrast of this feature is ~ 30 mK, and for even modest Δz , it becomes smooth and more easily subtracted by the foreground modes. In contrast, in the period prior to reionization, the 21 cm is thought to go into absorption at ~ 100 mK (Pritchard & Loeb 2010). Here, the spin temperature is strongly radiatively coupled to the cold gas. However, while synchrotron radiation contaminating $z \sim 10$ 21 cm radiation around the time of reionization is ~ 500 K (averaged across the sky), the contamination of an absorption signature at $z \sim 20$ is ~ 3000 K.

While the synchrotron emission from higher redshift 21 cm emission may be greater, the absorption signature has the additional feature that it is two-sided, both falling and rising. This generally makes foreground modes less parallel to the

signal and leads to less signal loss. For our purposes here, we can again develop a simple model where T_b is a Gaussian function parameterized by a central redshift and width σ_z . Fig. 12 is analogous to Fig. 11 except that it considers this era. Generally much more cosmological signal remains orthogonal to the foreground modes.

7.7. Susceptibility to Faraday rotation

In addition to the Stokes-I component of the foregrounds, synchrotron emission is also polarized. The polarized emission is subject to Faraday rotation as it traverses various column densities of free electrons in magnetic fields. If a receiver were designed to be purely sensitive to Stokes-I, then this spectral fluctuation could be ignored. Generally though, a receiver will have some level of response to polarized signals, and Faraday rotation produces a signal which oscillates in frequency. The rotation angle is $(\text{RM}) \cdot \lambda^2$, where RM is the rotation measure. Even for modest RM, the rotation angle can vary rapidly over our band. This is especially problematic for the global 21 cm signal because the rotation measure varies over the sky, increasing the number of contaminated modes in the data.

Extragalactic sources are subject to Faraday rotation through the screen of the entire Milky Way. This has been measured recently by Oppermann et al. (2012, 2014), and is shown in Fig. 13. Each point source will generally have a different rotation measure and require a new degree of freedom to describe. Bernardi et al. (2013) has conducted deep, wide-field observations of polarization of galactic synchrotron emission at 189 MHz. They find only one source out of 70 with $S > 4$ Jy with polarization at the level of 1.8% in 2400 deg^2 , with all others falling below 2% polarization fraction.

Polarization and Faraday rotation of galactic synchrotron emission is more complex because synchrotron emission is interspersed with the Faraday screen. Bernardi et al. (2013) find peak polarized emission at ~ 13 K and rotation measures mostly below 10 rad m^{-2} . The low rotation measures and polarization fraction are thought to be due to a depolarization horizon (Landecker et al. 2010; Uyaniker et al. 2003) within the galaxy, a distance beyond which most of the emission is depolarized along the line of sight.

The global 21 cm experiments to-date have been single-element, making it difficult to make a precise comparison with interferometric galaxy polarization surveys conducted at similar frequencies (Pen et al. 2009; Bernardi et al. 2013). Vinyaikin et al. (1996) measured 3.5 ± 1.0 K at 88 MHz (24° FWHM, phased array) and 2.15 ± 0.25 K at 200 MHz (8° FWHM, single-dish) for the value of $\sqrt{Q^2 + U^2}$ toward the North Celestial Pole.

To assess the galactic polarization field, we use the model of Waelkens et al. (2009), which self-consistently simulates the polarized emission and rotation measure. Fig. 13 shows the eigenvalue spectrum of the Stokes-Q spectral modes. Unlike Stokes-I, polarization is spectrally much more complex, and its covariance is spread among many modes. Ideally, all of these modes would be suppressed well below the noise level of the experiment so they do not need to be estimated and subtracted. Based on the measured amplitudes of the polarized signals above, suppose that polarization is 10^{-3} of intensity. To achieve a 10^5 suppression, the polarization leakage must be kept below 1%.




Figure 13. *Top:* Galactic rotation measure map of Oppermann et al. (2014), in units of rad m^{-2} , with color range truncated at $\pm 100 \text{ rad m}^{-2}$ to show more structure out of the plane. *Bottom:* The eigenvalue spectrum of Stokes Q emission in the galactic model of Waelkens et al. (2009), normalized to 1. This component falls off very slowly because the rotation measure introduces many degrees of freedom that vary across the sky.

8. CONCLUSION

Measurements of the global 21 cm signal are very challenging due to 1) astrophysical foregrounds and their interaction with instrumental systematics, 2) terrestrial radio interference, and 3) the ionosphere. Here we have examined the first issue, especially in regard to instrumental response. Following $z \sim 1$ 21 cm literature (Switzer et al. 2013; Masui et al. 2013; Chang et al. 2010), we develop a new method where foregrounds can be jointly estimated with the monopole spectrum. This relies on the fact that foregrounds vary across the sky, while the cosmological signal is constant. This idea also extends L13, who argue that surveys with moderate angular

resolution covering much of the sky are able to better discriminate between the monopole 21 cm signal and foregrounds.

The key observation arising from our cleaning method is that the instrument should be designed to minimize the generation of new spectral degrees of freedom from the foregrounds. For example, if each line of sight has a slightly different passband calibration, it also requires a new spectral degree of freedom to describe the foregrounds there. In this sense, the instrument “spreads out” the variance over modes that ultimately require more aggressive cleaning and signal loss. In contrast, a constant passband calibration error does not increase the rank of the foreground spectral covariance, and its effect is primarily aesthetic. This is fortuitous because obtaining a smooth spectral calibration of an instrument at these frequencies would require very large, expensive structures that are black in radio wavelengths. Simply requiring that galactic foregrounds be smooth to a few percent provides a sufficient passband calibration for signal recovery. In contrast, significant effort must be put into maintaining passband stability to at least $\sim 10^{-4}$. This may require using much of the instrument’s sensitivity to integrate against a stable reference.

Polarization to intensity leakage is another example of an instrumental systematic that increases the rank of the foreground covariance. It allows spectral oscillations in the Faraday-rotating polarization signal to contaminate the spectral intensity measurement. Because of depolarization, sky polarization is only a fraction of the intensity, and our estimates of the instrumental constraints are more lax, $< 10^{-2}$.

In attempting to measure the global 21 cm signal, one is faced with bright contaminating foregrounds that can interact with instrumental systematics in non-trivial ways, overwhelming the cosmological signal that one seeks to measure. In this paper, we have developed methods that bear similarities to various previously-proposed intuitive data analysis techniques. However, our methods arise from a rigorous, self-consistent framework that allows unknown and unanticipated foreground properties to be derived from real data. Such a framework also provides guidance for instrument design. If design requirements can be adequately met, high-significance measurements of the global 21 cm signal will be possible, providing direct access to the rich and complex astrophysics of the first luminous objects and reionization.

ES thanks Ue-Li Pen for conversations and feedback, and for stimulating the approach in collaboration on $z \sim 1$ 21 cm data analysis from the Green Bank Telescope. AL would like to thank Matt McQuinn for conversations. We would like to thank the Ohio 21 cm workshop for fostering conversations that initiated the project. AL acknowledges support from NSF grants AST-0804508, AST-1129258, AST-1125558.

APPENDIX

A. RELATION BETWEEN COVARIANCE ADJUSTMENT AND THE MAXIMUM-LIKELIHOOD LINEAR ESTIMATOR

In Sec. 3.2, we claim that the estimator $\mathbf{m} = (\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} \mathbf{x}$ is equivalent to the projection operation, $\mathbf{m} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T (1 - \Pi) \mathbf{y}$ where $\Pi = \Sigma \mathbf{Z} (\mathbf{Z}^T \Sigma \mathbf{Z})^{-1} \mathbf{Z}^T$. The two essential conditions here are that $\mathbf{X}^T \mathbf{Z} = 0$ (that the signal and non-signal vectors are orthogonal) and that the vectors in \mathbf{X} and \mathbf{Z} together span the space. The operations \mathbf{X}^T and \mathbf{Z}^T can be understood as projecting onto the signal basis and the everything-but-signal basis, respectively. This proof follows Lemma 2b of Rao (1967). Writing out the terms, we would like to prove

$$(\mathbf{X}^T \Sigma^{-1} \mathbf{X})^{-1} \mathbf{X}^T \Sigma^{-1} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T - (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \Sigma \mathbf{Z} (\mathbf{Z}^T \Sigma \mathbf{Z})^{-1} \mathbf{Z}^T. \quad (\text{A1})$$

Multiplying on the right by \mathbf{X} is trivially true because $\mathbf{Z}^T \mathbf{X} = 0$ and the other terms are the identity matrix. This checks just the \mathbf{X} subspace of the equality. To prove general equality we need to check the \mathbf{Z} subspace. Multiply on the right by \mathbf{Z} and from the left by $\mathbf{X}^T \Sigma^{-1} \mathbf{X}$ (which will allow simplification)

$$\mathbf{X}^T \Sigma^{-1} \mathbf{Z} = -(\mathbf{X}^T \Sigma^{-1} \mathbf{X})(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \Sigma \mathbf{Z} (\mathbf{Z}^T \Sigma \mathbf{Z})^{-1} \mathbf{Z}^T \mathbf{Z} \quad (\text{A2})$$

$$\mathbf{X}^T \Sigma^{-1} \mathbf{Z} = -(\mathbf{X}^T \Sigma^{-1} \mathbf{X})(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \Sigma \mathbf{Z} [(\mathbf{Z}^T \mathbf{Z})^{-1} (\mathbf{Z}^T \Sigma \mathbf{Z})]^{-1} \quad (\text{A3})$$

$$\mathbf{X}^T \Sigma^{-1} \mathbf{Z} [(\mathbf{Z}^T \mathbf{Z})^{-1} (\mathbf{Z}^T \Sigma \mathbf{Z})] = -(\mathbf{X}^T \Sigma^{-1} \mathbf{X})(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \Sigma \mathbf{Z} \quad (\text{A4})$$

$$\mathbf{X}^T \Sigma^{-1} [\mathbf{Z} (\mathbf{Z}^T \mathbf{Z})^{-1} \mathbf{Z}^T + \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T] \Sigma \mathbf{Z} = 0. \quad (\text{A5})$$

Because \mathbf{X} and \mathbf{Z} span the space, $\mathbf{Z} (\mathbf{Z}^T \mathbf{Z})^{-1} \mathbf{Z}^T + \mathbf{X} (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T = \mathbf{1}$, then noting that $\mathbf{X}^T \mathbf{Z} = 0$ proves the identity. In many of the methods in the paper, the projection $\mathbf{\Pi}$ is reduced to $(1 - \mathbf{F} \mathbf{F}^T)$, completely projecting out a set of modes. The modes in \mathbf{F} clearly cannot span the space. Further, we let the modes \mathbf{F} overlap with the signal. These choices make the estimator less formally optimal, but more robust to very bright, non-Gaussian foregrounds.

B. EXTENDING THE MONOPOLE CONSTRAINT: THE GROWTH MODEL

The separable model for the monopole can be extended to a more general spatial-spectral template scheme. Rather than a simple signal model $\alpha \mathbf{x} \mathbf{e}_0^T$, the mean of the observed data can be described by the outer products spectral modes \mathbf{u}_i and spatial modes \mathbf{v}_i (not necessarily mutually orthogonal). In terms of the $N_\nu \times N_{\text{pix}}$ data cube \mathbf{Y} ,

$$\mathbf{Y} = \mathbf{U} \mathbf{X} \mathbf{V}^T = \sum_{i,j} x_{ij} \mathbf{u}_i \mathbf{v}_j^T. \quad (\text{B1})$$

This extended model could represent spatial and spectral templates of galactic emission in addition to the monopole signal. The noise terms in the model would then describe any residuals with respect to this model, written as

$$\mathbf{Y} = \mathbf{U} \mathbf{X} \mathbf{V}^T + \Sigma^{1/2} \mathbf{E} \Phi^{1/2} \quad (\text{B2})$$

This separable estimation problem is considered in detail by Kollo & von Rosen (2005). (To reach their notation, right multiply by $\Phi^{-1/2}$.) Form the weighted (ν, ν') sample covariance with the spatial components of the signal projected out, as

$$\mathbf{S} = \mathbf{Y} \Phi^{-1} (\mathbf{1} - \mathbf{V} \mathbf{W} \mathbf{V}^T) \mathbf{Y}^T \quad (\text{B3})$$

where $\mathbf{W}_\mathbf{V} = (\mathbf{V}^T \Phi^{-1} \mathbf{V})^{-1} \mathbf{V}^T \Phi^{-1}$ is the standard maximum-likelihood estimator for the amplitudes of the spatial vectors in \mathbf{V} . The maximum likelihood estimator for \mathbf{X} and full-rank spectral foregrounds \mathbf{S} is

$$\hat{\mathbf{X}} = (\mathbf{U}^T \mathbf{S}^{-1} \mathbf{U})^{-1} \mathbf{U}^T \mathbf{S}^{-1} \mathbf{Y} \mathbf{W}_\mathbf{V}^T. \quad (\text{B4})$$

Kollo & von Rosen (2005) describe extensions of this more general estimation problem to finite rank \mathbf{S} .

REFERENCES

- Barkana, R., & Loeb, A. 2001, *Phys. Rep.*, 349, 125
 Bennett, C. L., Larson, D., Weiland, J. L., et al. 2013, *ApJS*, 208, 20
 Bernardi, G., McQuinn, M., & Greenhill, L. J. 2014, *ArXiv e-prints*, arXiv:1404.0887
 Bernardi, G., Greenhill, L. J., Mitchell, D. A., et al. 2013, *ApJ*, 771, 105
 Bowman, J. D., & Rogers, A. E. E. 2010, *Nature*, 468, 796
 Bromm, V. 2013, *Reports on Progress in Physics*, 76, 112901
 Burns, J. O., Lazio, J., Bale, S., et al. 2012, *Advances in Space Research*, 49, 433
 Chang, T.-C., Pen, U.-L., Bandura, K., & Peterson, J. B. 2010, *Nature*, 466, 463
 de Oliveira-Costa, A., Tegmark, M., Gaensler, B. M., et al. 2008, *MNRAS*, 388, 247
 Di Matteo, T., Perna, R., Abel, T., & Rees, M. J. 2002, *ApJ*, 564, 576
 Dillon, J. S., Liu, A., Williams, C. L., et al. 2014, *PRD*, 89, 023002
 Fan, X., Strauss, M. A., Becker, R. H., et al. 2006, *AJ*, 132, 117
 Field, G. B. 1958, *Proc. IRE*, 46, 240
 Furlanetto, S. R., Oh, S. P., & Briggs, F. H. 2006, *Phys. Rep.*, 433, 181
 Gleser, L. J., & Olkin, I. 1972, in *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Theory of Statistics*,
 Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Theory of Statistics (Berkeley, Calif.: University of
 California Press), 541–568
 Górski, K. M., Hivon, E., Banday, A. J., et al. 2005, *ApJ*, 622, 759
 Knox, L., Scoccimarro, R., & Dodelson, S. 1998, *PRL*, 81, 2004
 Kollo, T., & von Rosen, D. 2005, *Advanced Multivariate Statistics with Matrices*
 Landecker, T. L., Reich, W., Reid, R. I., et al. 2010, *A&A*, 520, A80
 Liu, A., Pritchard, J. R., Tegmark, M., & Loeb, A. 2013, *PRD*, 87, 043002
 Liu, A., & Tegmark, M. 2012, *MNRAS*, 419, 3491
 Masui, K. W., Switzer, E. R., Banavar, N., et al. 2013, *ApJ*, 763, L20
 Morales, M. F., Hazelton, B., Sullivan, I., & Beardsley, A. 2012, *ApJ*, 752, 137
 Oppermann, N., Junklewitz, H., Robbers, G., et al. 2012, *A&A*, 542, A93
 Oppermann, N., Junklewitz, H., Greiner, M., et al. 2014, *ArXiv e-prints*, arXiv:1404.3701
 Paciga, G., Albert, J. G., Bandura, K., et al. 2013, *MNRAS*, 433, 639
 Parsons, A. R., Liu, A., Aguirre, J. E., et al. 2013, *ArXiv e-prints*, arXiv:1304.4991
 Pen, U.-L., Chang, T.-C., Hirata, C. M., et al. 2009, *MNRAS*, 399, 181
 Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2013, *ArXiv e-prints*, arXiv:1303.5076

- Pober, J. C., Parsons, A. R., Aguirre, J. E., et al. 2013, *ApJL*, 768, L36
- Pritchard, J. R., & Loeb, A. 2008, *PRD*, 78, 103511
- . 2010, *PRD*, 82, 023006
- Rao, C. R. 1967, in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics (Berkeley, Calif.: University of California Press), 355–372
- Switzer, E. R., Masui, K. W., Bandura, K., et al. 2013, *MNRAS*, 434, L46
- Tegmark, M., de Oliveira-Costa, A., & Hamilton, A. J. 2003, *Phys. Rev. D*, 68, 123523
- Uyaniker, B., Landecker, T. L., Gray, A. D., & Kothes, R. 2003, *ApJ*, 585, 785
- van Haarlem, M. P., Wise, M. W., Gunst, A. W., et al. 2013, *A&A*, 556, A2
- Vedantham, H. K., Koopmans, L. V. E., de Bruyn, A. G., et al. 2014, *MNRAS*, 437, 1056
- Vinyaikin, E. N., Kuznetsova, I. P., Paseka, A. M., Razin, V. A., & Teplykh, A. I. 1996, *Astronomy Letters*, 22, 582
- Voytek, T. C., Natarajan, A., Jáuregui García, J. M., Peterson, J. B., & López-Cruz, O. 2014, *ApJL*, 782, L9
- Waelkens, A., Jaffe, T., Reinecke, M., Kitauro, F. S., & Enßlin, T. A. 2009, *A&A*, 495, 697
- Wouthuysen, S. A. 1952, *Physica*, 18, 75
- Zahn, O., Reichardt, C. L., Shaw, L., et al. 2012, *ApJ*, 756, 65